

IBM Spectrum Copy Data Management[™] 2.2.12.0

User's Guide

Table of Contents

- Audience and Purpose 9**
- IBM Spectrum Copy Data Management Overview10**
- IBM Spectrum Copy Data Management Workflow 11**
- User Administration and Security Management14**
- Installation and Setup18**
 - Deployment Checklist19
 - Access and Default Credentials20
 - System Requirements 23
 - File System Requirements28
 - InterSystems Caché Requirements34
 - Oracle Requirements 41
 - SAP HANA Requirements51
 - Microsoft SQL Server Requirements57
 - AWS Requirements 62
 - Install IBM Spectrum Copy Data Management as a Virtual Appliance63
 - Start IBM Spectrum Copy Data Management65
- Dashboard67**
- Sites & Providers69**
 - Sites & Providers Overview70
 - Add a Site72
 - Edit a Site73
 - Delete a Site74
 - Register a Provider75
 - View a Provider95
 - Edit a Provider98
 - Unregister a Provider100
 - Add Credentials to a Virtual Machine101
- Role-Based Access Control103**

Role-Based Access Control Overview	104
Configure Resource Pools	106
Configure Roles	109
Configure Accounts	111
VMware Admin Role-Based Access Control Configuration	113
NetApp ONTAP Admin Role-Based Access Control Configuration	115
IBM Admin Role-Based Access Control Configuration	117
HPE Nimble Storage Admin Role-Based Access Control Configuration	118
Pure Storage FlashArray Admin Role-Based Access Control Configuration	120
Configure Tenants	122
Best Practices for Configuring Tenants	124
Identities	126
Identities Overview	127
Add a Key	128
Add a Credential	132
Configure SLA Policies	133
Configure Scripts	152
Schedules	154
Create a Schedule	155
Edit a Schedule	157
Delete a Schedule	159
Jobs	160
Jobs Overview	161
Start, Pause, and Hold a Job Session	162
Monitor a Job Session	164
Job Definition Overview	166
Inventory Jobs	169
Create an Inventory Job Definition - Database	170
Create an Inventory Job Definition - File System	173
Create an Inventory Job Definition - DellEMC Unity	175
Create an Inventory Job Definition - IBM Spectrum Accelerate	177

Create an Inventory Job Definition - IBM Spectrum Protect Snapshot 179

Create an Inventory Job Definition - IBM Spectrum Virtualize 181

Create an Inventory Job Definition - NetApp ONTAP Storage183

Create an Inventory Job Definition - NetApp ONTAP File 186

Create an Inventory Job Definition - HPE Nimble Storage 189

Create an Inventory Job Definition - Pure Storage FlashArray 192

Create an Inventory Job Definition - VMware 194

Backup Jobs196

 Create a Backup Job Definition - InterSystems Caché 197

 Create a Backup Job Definition - SAP HANA200

 Create a Backup Job Definition - Oracle 204

 Create a Backup Job Definition - SQL211

 Create a Backup Job Definition - File System216

 Create a Backup Job Definition - DellEMC Unity220

 Create a Backup Job Definition - IBM Spectrum Accelerate 223

 Create a Backup Job Definition - IBM Spectrum Virtualize 227

 Create a Backup Job Definition - NetApp ONTAP 231

 Create a Backup Job Definition - HPE Nimble Storage 234

 Create a Backup Job Definition - Pure Storage FlashArray 238

 Create a Backup Job Definition - VMware242

 Create VMware Backup Job Proxies 248

Restore Jobs251

 Create a Restore Job Definition - InterSystems Caché 252

 Create a Restore Job Definition - SAP HANA 257

 Create a Restore Job Definition - Oracle264

 Create a Restore Job Definition - SQL 270

 Create a Restore Job Definition - File System 277

 Create a Restore Job Definition - DellEMC Unity 281

 Create a Restore Job Definition - IBM Spectrum Accelerate287

 Create a Restore Job Definition - IBM Spectrum Virtualize293

 Create a Restore Job Definition - NetApp ONTAP299

Create a Restore Job Definition - HPE Nimble Storage	306
Create a Restore Job Definition - Pure Storage FlashArray	312
Create a Restore Job Definition - VMware	317
Restore Jobs - Rename Mount Points and Initialization Parameter Options	331
Using State and Status Arguments in Postscripts	334
System Jobs	335
Maintenance Job	336
Create a Report Job Definition	337
Create a Script Job Definition	339
Edit a Job Definition	342
Delete a Job Definition	343
Search	344
Search Overview	345
Search for Objects	346
View Object Details	351
View NetApp ONTAP File Details	352
Find and Restore a File	354
Download Search Results	356
Browse Inventory	357
Report	359
Report Overview	360
Run a Report	361
Create a Customized Report	363
Edit a Customized Report	364
Download a Report	365
Delete a Generated Report	367
Application Reports	368
Application Configuration Report	369
Application RPO Compliance Report	372
System Management Reports	376
Catalog Summary Report	377

Configuration Report	379
Job Report	383
Job Sessions Report	386
File Analytics Reports	388
File Usage by Owner Report	390
Files By Age Report	392
Files By Category Report	394
Files By Size Report	396
Protection Compliance Reports	398
DellEMC Unity RPO Compliance Report	400
File System RPO Compliance Report	404
HPE Nimble Storage RPO Compliance Report	407
IBM Spectrum Accelerate RPO Compliance Report	411
IBM Spectrum Virtualize RPO Compliance Report	413
NetApp ONTAP Protection Usage Report	417
NetApp ONTAP RPO Compliance Report	420
Pure Storage FlashArray RPO Compliance Report	424
Recovery Points Report	427
Unprotected Virtual Machines Report	429
VMware RPO Compliance Report	431
Storage Protection Reports	435
NetApp ONTAP OSSV Relationship Status Report	437
NetApp ONTAP Overprotected Volumes Report	439
NetApp ONTAP Qtree Protection Status Report	441
NetApp ONTAP Underprotected Volumes Report	443
NetApp ONTAP Volume Protection Status Report	445
NetApp ONTAP Transition Dependency Report	447
Storage Utilization Reports	449
DellEMC Unity File Systems Report	453
DellEMC Unity LUNs Report	455
DellEMC Unity Pools Report	458

HPE Nimble Storage Volumes Report	460
IBM Spectrum Accelerate Pools Report	463
IBM Spectrum Accelerate Volumes Report	465
IBM Spectrum Virtualize Consistency Groups Report	467
IBM Spectrum Virtualize Pools Report	469
IBM Spectrum Virtualize Volumes Report	472
Instant Disk Restore Volumes Report	475
NetApp ONTAP Aggregates Report	476
NetApp ONTAP LUNs Report	479
NetApp ONTAP Orphaned LUNs Report	482
NetApp ONTAP Quotas Report	484
NetApp ONTAP Snapshots Report	486
NetApp ONTAP Volumes Report	488
Pure Storage FlashArray Volumes Report	491
Storage Capacity Report	494
VM and Storage Mapping Report	496
VMware Datastores Report	498
VMware LUNs Report	501
VMware Orphaned Datastores Report	503
VMware Orphaned LUNs Report	505
VMware VM Snapshot Sprawl Report	507
VMware VM Sprawl Report	509
VMware VM Storage Report	512
Maintenance	513
Maintenance Overview	514
Log On to the Virtual Appliance	515
Set Time Zone	516
Collect Logs For Troubleshooting	518
Modifying Job Log Options	521
Manage the Administrative Console	523
Update IBM Spectrum Copy Data Management	524

To update your IBM Spectrum Copy Data Management appliance: 524

Manually Update Administrative Console RPMs When Upgrading from IBM Spectrum Copy Data Management 2.2.6 to 2.2.7 525

To update your IBM Spectrum Copy Data Management appliance through RPM update files: 525

Install the Marketplace RPM 527

Backup and Restore the Catalog 528

Modifying Network Settings 530

Upload an SSL Certificate 532

Horizontal Scale Out 534

Restoring A Snapshot From A FlexGroup Volume To Another FlexGroup Volume 539

Documentation and Support 541

Documentation Roadmap 542

About the Help System 543

Reference Topics 544

Virtual Machine Privileges 545

Search and Filter Guidelines 548

Select, Sort, and Reorder Columns 551

LDAP User Name Syntax 552

Return Code Reference 553

Frequently Asked Questions 556

Oracle Database Support FAQ 563

Microsoft SQL Server Support FAQ 584

Acronyms 598

Terminology 600

Terminology Changes 605

Trademarks 606

Index 608

Audience and Purpose

This PDF is intended for IBM Spectrum Copy Data Management users, system administrators, and the Super User. It contains information, procedures, and tips for the most commonly used functions.

System administrators can use this guide to help install, maintain, and start the application, manage users, and catalog resource information. Users can find procedures on how to search and browse for objects, generate and interpret reports, schedule jobs, and orchestrate Backup and Restore jobs.

IBM Spectrum Copy Data Management Overview

Organizations of all sizes need to modernize their IT processes to enable critical new use cases such as operational automation, DevOps and integration of system-of-record data with Cloud compute. They are equally challenged with improving management efficiencies for long established IT processes such as data protection, disaster recovery, reporting and test and development.

IBM Spectrum Copy Data Management is a Copy Data Management (CDM) platform that can bring modernization to an existing environment without disruption. IBM Spectrum Copy Data Management delivers “in-place” copy data management to enterprise storage arrays from EMC, HPE Nimble Storage¹, IBM, NetApp and Pure Storage, allowing the IT team to make use of its existing infrastructure and data in a manner that is efficient, automated, scalable and easy to use. IBM Spectrum Copy Data Management modernizes IT processes, enables key use cases, and does it all without additional hardware.

¹Replication is not supported on HPE Nimble Storage.

IBM Spectrum Copy Data Management Workflow

The following is a typical IBM Spectrum Copy Data Management workflow including starting IBM Spectrum Copy Data Management, registering a provider, cataloging data, searching for objects, generating reports, and copying and using data:

Start IBM Spectrum Copy Data Management

1. From a supported browser, enter the following URL:

`https://<HOSTNAME>:8443/portal/`

where <HOSTNAME> is the IP address of the virtual machine where the application is deployed. This connects you to IBM Spectrum Copy Data Management.

2. Enter your user name and password, which is provided by the IBM Spectrum Copy Data Management System Administrator. If this is your first time logging on to IBM Spectrum Copy Data Management as a Super User, the default user name is **admin** and the default password is **password**. You are prompted to reset the default Super User password.
3. Click **Sign In**. The application launches.

Register a Provider

Add providers such as application servers, DellEMC, IBM, or NetApp storage devices, or VMware ESX resources to the Inventory by registering them.

1. Click the **Configure**  tab and the **Sites & Providers**  view.
2. In the Provider Browser pane, right-click a provider category and then click **Register** . The Register dialog opens.
3. Populate the fields in the dialog, including name, host address, port, user name, and password.

Create an Inventory Job Definition

An Inventory job definition provides the framework to collect and catalog information about objects on a registered provider.

1. Click the **Jobs**  tab and the **Policies**  view.
2. Select a provider type, click **New**  and then select **Inventory**.
3. Select one or more resources to catalog from the list of available providers.
4. Select the options for your job. Also, enter notifications. If notification options are enabled, an email

message with information about the status of each task is sent when the job completes.

5. Optionally, select one or more defined schedules for your job and save the job definition.

Run a Job

A job that is based on an Inventory job definition discovers object information, catalogs it, and populates the IBM Spectrum Copy Data Management database.

1. Click the **Jobs**  tab
2. Select the job to run by clicking in the row containing the job name.
3. From the Operations drop-down menu, select **Start** , or right-click the job name and select **Start** .

Search for Objects

Search for objects on specified cataloged nodes. Use advanced search filters to tailor the search.

1. Click the **Search**  tab.
2. Enter the object name or character string to search on.
3. Click **Search Now**. The list of objects displays.
4. Optionally, click **Advanced Search** and apply filters such as catalog, object type, location, and name.
5. Click **Search** . The list of objects that meet all the criteria displays.

Generate Reports

Run a report to summarize information about cataloged nodes as well as the data and resources that reside on them.

1. Click the **Report**  tab and the **Reports**  view.
2. Select one of the predefined reports to run by clicking in the row containing the job name.
3. Click **Run**  to run the report using default parameters.
4. Optionally, select a predefined report from the **Report Browser** pane and select report parameter values in the **Parameters** pane.
5. Click **Run** . The customized report data is returned in the **Report** pane.

Backup

Create copies of your data. The RPO and copy data parameters are defined in an SLA Policy, which is then applied to the Backup job definition along with a specified activation time to meet your copy data criteria.

1. Click the **Jobs**  tab.
2. Select a provider type, click **New**  and then select **Backup**.

3. Select providers to copy or protect, as well as a storage workflow that meets your copy data criteria.
4. Complete the job definition including notification and other options. Save the job definition.
5. Run the backup job. The selected source data is copied to the destination in accordance with the defined job parameters.

Restore

Leverage Copy Data Management technology for testing, cloning, and recovering copy data.

1. Click the **Jobs**  tab.
2. Select a provider type, click **New**  and then select **Restore** .
3. Select a workflow, providers to reuse or recover, as well as destinations
4. Complete the job definition including notification and other options. Save the job definition.
5. Run the restore job. The selected data is made available for use in accordance with the defined job parameters.

RELATED TOPICS:

- [Start IBM Spectrum Copy Data Management](#) on page **65**
- [Sites & Providers Overview](#) on page **70**
- [Start, Pause, and Hold a Job Session](#) on page **162**
- [Search for Objects](#) on page **346**
- [Report Overview](#) on page **360**

User Administration and Security Management

IBM Spectrum Copy Data Management provides users the opportunity to rapidly locate files and objects on DellEMC, IBM, NetApp, and Pure Storage FlashArray devices along with VMware ESXi, Oracle, and SQL hosts. IBM Spectrum Copy Data Management then stores this information so you can report on it. The reports provide a basis for users to take administrative actions towards efficient management of the DellEMC, IBM, NetApp, and/or Pure Storage FlashArray storage devices, along with VMware, Oracle, and SQL hosts and resources.

IBM Spectrum Copy Data Management security objectives are:

- Identify and authenticate users prior to providing any of its services.
- Ensure all functions are authorized.
- Protect confidentiality of DellEMC, IBM, NetApp, Pure Storage FlashArray, VMWare, Oracle, and SQL server credentials by encrypting them when stored and in transit.
- Prevent bypass of and tampering with its security functions through perimeter hardening and use of secure transmission protocols.

Note that IBM Spectrum Copy Data Management uses FIPS compliant encryption algorithms.

Identification and Authentication

All services require some form of authentication.

Users are uniquely identified by entering a user name and password. System Administrators have the option of adding native users or importing groups of provisioned users through LDAP authentication. Native user names are not case sensitive. LDAP user name case sensitivity relies on the configuration of your LDAP server.

User Data Security

IBM Spectrum Copy Data Management employs role-based access control to provisioned users:

- Native users or members of imported LDAP groups are assigned to roles.
- Roles contain collections of permissions that allow access to IBM Spectrum Copy Data Management functionality.

Sensitive data is encrypted when stored.

Data in transit is also protected. IBM Spectrum Copy Data Management protects the confidentiality of the user and system credentials. Sensitive data is encrypted or transported using SSL and HTTPS. The user login is protected via HTTPS for browser client to IBM Spectrum Copy Data Management server login, and via LDAP/S for communication with the LDAP directory server. For backend processes, protection is secured via HTTPS authentication to the storage system and ESXi.

IBM Spectrum Copy Data Management identifies the following types of sensitive data: native user credentials, DellEMC, IBM, NetApp, and Pure Storage FlashArray storage system credentials, VMware/ESX host credentials, and user credentials.

Security Management

Security management identifies the interfaces that manage the security functions in the IBM Spectrum Copy Data Management application. Only an authenticated, authorized user can configure the security functions. Examples of security management include adding users, assigning roles, configuring IBM Spectrum Copy Data Management to use LDAP, and configuring IBM Spectrum Copy Data Management to use HTTPS. Following are the security management functions in IBM Spectrum Copy Data Management:

- Adding, editing, and deleting a user
- Assigning roles to a user
- Configuring authentication mode
- Configuring LDAP
- Importing certifications
- Configuring HTTPS

Management and Operation Functions

Management and operation functions include session timeout, log on credentials, and role-based access control mechanism:

- The session timeout specifies the time-out period assigned for the application in minutes. If the user does not refresh or request a window within the time-out period, the session ends automatically. Session timeout is set for 30 minutes and cannot be changed.
- Users are uniquely identified by entering a user name and password.
- Role-based access control is employed. Once a user is added to IBM Spectrum Copy Data Management, either as a native user or imported as part of an LDAP group, the user is assigned to specified resource pools and roles.

Encryption

IBM Spectrum Copy Data Management provides encryption solutions for complete security. The solution includes certificates, use of HTTPS, and safe storage of passwords in the database. Sensitive data such as data in transit is encrypted or transported using SSL and HTTPS. User credentials such as passwords are safely stored in the IBM Spectrum Copy Data Management database. Obtaining and storing this sensitive data constitutes the basic function of the IBM Spectrum Copy Data Management application. This data is subject to the user data security requirements.

Ports

The following ports are used by IBM Spectrum Copy Data Management:

Ports

Port	Service	Comment
22	OpenSSH 5.3 (protocol 2.0)	Port open within the firewall.
25	smtp, non-SSL connection for Simple Mail Transfer Protocol	Service used by IBM Spectrum Copy Data Management
68	bootpc in DHCP clients, DHCP Listener UDP	
80/443	http/https	Service used by IBM Spectrum Copy Data Management
389	LDAP, non-SSL connection for Lightweight Directory Access Protocol	Service used by IBM Spectrum Copy Data Management
443	smtp, SSL connection for Simple Mail Transfer Protocol	Service used by IBM Spectrum Copy Data Management
636	LDAP, SSL connection for Lightweight Directory Access Protocol	Service used by IBM Spectrum Copy Data Management
1433	sql, SQL Service	Service used by IBM Spectrum Copy Data Management
4369	epmd, Erlang port mapper	Service used by IBM Spectrum Copy Data Management
5432	postgresql, PostgreSQL DB 8.4.1-8.4.4	Service used by IBM Spectrum Copy Data Management
5480	ssl/http, vami	Port open within the firewall, IBM Spectrum Copy Data Management versions 2.2.5 and earlier
5985	WinRM, Windows Remote Management	Service used by IBM Spectrum Copy Data Management
8090	adminconsole, IBM Spectrum Copy Data Management Administrative Console	Port open within the firewall
8092	adminconsole, IBM Spectrum Copy Data Management Administrative Console	Port open within the firewall, IBM Spectrum Copy Data Management version 2.2.6 only

Port	Service	Comment
8443	ssl/http, Apache Tomcat/Coyote JSP engine 1.1	Port open within the firewall
8761	Discovery Server	Service used by IBM Spectrum Copy Data Management. Locates registered micro services.
27017	MongoDB mongod	Service used by IBM Spectrum Copy Data Management
27018	MongoDB mongod	Service used by IBM Spectrum Copy Data Management
55672	rabbitMQ, RabbitMQ administrative	Service used by IBM Spectrum Copy Data Management

RELATED TOPICS:

- [Register a Provider](#) on page **75**
- [Role-Based Access Control Overview](#) on page **104**

Installation and Setup

The topics in the following section cover installing IBM Spectrum Copy Data Management and system requirements.

Deployment Checklist

Following are the pre-deployment, deployment, and post-deployment procedures. This checklist is for IBM Spectrum Copy Data Management deployment to a VMware appliance host.

Pre-Deployment Checklist

Step	Action	Related Topic	✓
•	Reference the locations of access and default credentials for IBM Spectrum Copy Data Management.	Access and Default Credentials on page 20	

Deployment Checklist

Step	Action	Related Topic	✓
1	Install IBM Spectrum Copy Data Management by deploying an OVF template to create a virtual appliance containing the application on a VMware host.	Install IBM Spectrum Copy Data Management as a Virtual Appliance on page 63	
2	To access IBM Spectrum Copy Data Management, upload a valid product key to the virtual machine where IBM Spectrum Copy Data Management is deployed.	Install IBM Spectrum Copy Data Management as a Virtual Appliance on page 63	
3	Launch IBM Spectrum Copy Data Management to set a new Super User password.	Start IBM Spectrum Copy Data Management on page 65	
4	Configure LDAP Authentication and establish a secure connection to the LDAP server.	Register a Provider on page 75	
5	Create accounts through role-based access control.	Role-Based Access Control Overview on page 104	

Post-Deployment Checklist

Step	Action	Related Topic	✓
•	Launch IBM Spectrum Copy Data Management to begin using the application and its features.	Start IBM Spectrum Copy Data Management on page 65	
•	Update IBM Spectrum Copy Data Management to keep it current with new features, enhancements, and upgrades.	Update IBM Spectrum Copy Data Management on page 524	
•	Optionally, log on to the virtual appliance's web-based management console to review the configuration of the virtual appliance. Available information includes system settings, network, proxy settings, and available updates.	Manage the Administrative Console on page 523	

Access and Default Credentials

The following user accounts are accessible from a fresh IBM Spectrum Copy Data Management installation or deployment. Change the password for each account after installation or deployment in your environment.

Interface	Location	Default Credentials	Notes
IBM Spectrum Copy Data Management User Interface	<a href="https://<HOSTNAME_OR_IP>:8443/portal/">https://<HOSTNAME_OR_IP>:8443/portal/	Username: <code>admin</code> Password: <code>password</code>	This is the main IBM Spectrum Copy Data Management User Interface and the default account associated with that interface. When logging in for the first time, you will be prompted to update the password to something more secure. This account may be used to create other IBM Spectrum Copy Data Management users.

<p>IBM Spectrum Copy Data Management Administrative Console</p>	<p>https://<HOSTNAME_OR_IP>:8090/</p>	<p>Username: administrator Password: ecxadLG235</p>	<p>This interface is used to make administrative changes to the IBM Spectrum Copy Data Management server. In the drop down menu, select System to use the administrator account. Otherwise, select IBM Spectrum Copy Data Management to enter a valid IBM Spectrum Copy Data Management account username and password.</p>
<p>IBM Spectrum Copy Data Management Command Line / Console</p>	<p>SSH or Console</p>	<p>Username: administrator Password: ecxadLG235</p>	<p>The administrator account may be used to access the IBM Spectrum Copy Data Management server through SSH or through console access. This account may also be used to access the Administrative Console (above). Note: This account does not have sudo access. If elevated privileges are required, use the root account.</p>

<p>IBM Spectrum Copy Data Management Command Line / Console</p>	<p>SSH or Console</p>	<p>Username: <code>root</code> Password: <code>ecxDP758</code></p>	<p>The root account may be used to make changes to the IBM Spectrum Copy Data Management server through SSH or through console access. Use this account when elevated privileges are required.</p>
<p>IBM Spectrum Copy Data Management Agent User</p>	<p>Application Server - SSH or Console</p>	<p>Suggested Username: <code>ecxagent</code></p>	<p>When adding a database or file system, you must create an agent user on that system. For more information, see the appropriate topic, "Sample Configuration of an IBM Spectrum Copy Data Management Agent User" for the file system or application. This account is only needed on the client systems, not on the IBM Spectrum Copy Data Management OVA.</p>

System Requirements

Ensure that you have the required system configuration and browser to deploy and run IBM Spectrum Copy Data Management.

Upgrades to new versions of IBM Spectrum Copy Data Management will continue to use the previous network configuration found at this location: `/opt/vmware/share/vami/vami_config_net`.

Virtual Machine Installation

IBM Spectrum Copy Data Management is installed as a virtual appliance. Before deploying to the host, ensure you have the following:

- The correct OVF template, which has an OVA extension, and is approximately 1.4 GB
- vSphere 6.0, 6.5, or 6.7
 - Note:** For later versions of vSphere, the vSphere Web Client may be required to deploy IBM Spectrum Copy Data Management appliances.
- Network information and VMware host information
- Either an available static IP address to use or access to DHCP

For initial deployment, configure your virtual appliance to meet the following recommended minimum requirements:

- 64-bit dual core machine
- 16 GB memory

The appliance has three virtual disks that total 400 GB storage:

- 50 GB for operating system and application, which includes 16 GB for the swap partition, 256 MB for the boot partition, and the remainder for the root partition
- 100 GB for configuration data related to jobs, events, and logs
- 250 GB for Inventory data

It is recommended that the IBM Spectrum Copy Data Management appliance, storage arrays, hypervisors and application servers in your environment use NTP. If the clocks on the various systems are significantly out of sync, you may experience errors during application registration, Inventory, Backup, or Restore jobs. For more information about identifying and resolving timer drift, see the following VMware knowledge base article: [Time in virtual machine drifts due to hardware timer drift](#).

Browser Support

Run IBM Spectrum Copy Data Management from a computer that has access to the installed virtual appliance.

IBM Spectrum Copy Data Management was tested and certified against the following web browsers. Note that newer versions may be supported.

- Internet Explorer 11
- Microsoft Edge 20.10240
- Firefox 49.0
- Chrome 53.0.27

If your resolution is below 1024 x 768, some items may not fit on the window. Pop-up windows must be enabled in your browser to access the Help system and some IBM Spectrum Copy Data Management operations.

Application Server Requirements

- [File System Requirements](#) on page 28
- [InterSystems Caché Requirements](#) on page 34
- [Oracle Requirements](#) on page 41
- [SAP HANA Requirements](#) on page 51
- [Microsoft SQL Server Requirements](#) on page 57

DellEMC Unity Requirements

- EMC Unity 300, 400, 500, 600 (All-Flash and Hybrid Flash)
- EMC UnityVSA
- EMC VNXe 1600 running software version 3.1.3 and later
- EMC VNXe 3200 running software version 3.1.1 and later

HPE Nimble Storage

- HPE Nimble Storage running version 5.2 or later.

Note: Firmware version 5.2 or later is required for iSCSI and Fibre Channel.

IBM Storage Requirements

- IBM storage systems running IBM Spectrum™ Virtualize Software version 7.3 and later/8.1.2 and later, including IBM SAN Volume Controller, IBM Storwize, and IBM FlashSystem V9000 and 9100 systems. Outdated SAN Volume Controller (SVC) firmware may result in a storage warmstart. SVC firmware should be updated to the latest supported version to ensure system reliability.
- HyperSwap is supported on both IBM Storwize V7000 and FS9100 for both physical and virtual server types for Volume, Virtual Machine, and SQL applications.
- IBM storage systems running IBM Spectrum Accelerate version 11.5.3 and later, including FlashSystem A9000/A9000R and IBM XIV storage systems

- IBM FlashCopy Manager systems running IBM Tivoli® Storage FlashCopy® Manager version 4.1.3 and later

Note: IBM providers must be registered by an IBM user with administrator-level privileges.

NetApp ONTAP Requirements

Any NetApp ONTAP storage system model running specifically listed Data ONTAP versions operating in 7-Mode is supported. The following versions of Data ONTAP are supported:

- Data ONTAP 8.1.0 and later operating in 7-Mode
- Data ONTAP 8.2.0 and later operating in 7-Mode
- Clustered Data ONTAP 8.1 and later
- Clustered Data ONTAP 8.2 and later
- Clustered Data ONTAP 8.3 and later
- Clustered Data ONTAP 9.4 and later
- Data ONTAP 9.x

Note that IBM Spectrum Copy Data Management does not support new features specifically introduced in ONTAP 9.x. IBM Spectrum Copy Data Management was tested against Data ONTAP 9.3.

IBM Spectrum Copy Data Management 2.2.7 and later supports NetApp MetroCluster configurations running ONTAP 9.x and later. After successful completion of MetroCluster Switchover or Switchback operations, mirror or vault relationships must be reestablished through an IBM Spectrum Copy Data Management job. See [Create a Backup Job Definition - NetApp ONTAP](#) on page 231.

Clustered Data ONTAP providers must be registered with a cluster administrator account. Cluster peering must be enabled. Peer relationships enable communication between SVMs. See NetApp ONTAP's [Cluster and Vserver Peering Express Guide](#).

Note: Ensure that TLS protocol is enabled on the NetApp storage system by setting the `tls.enable` option to ON. For TLS to take effect on HTTPS, ensure that the `httpd.admin.ssl.enable` option is also set to ON. See [Enabling or disabling TLS](#) on NetApp's Support site.

NetApp ONTAP File Inventory Job Requirements:

IBM Spectrum Copy Data Management uses SnapDiff in the NetApp ONTAP file level jobs to perform cataloging based on snapshot differences. SnapDiff cataloging is supported on storage system models running the following versions of Data ONTAP:

The following options must be enabled on the volume of the NetApp storage system to catalog:

- **create_ucose** and **convert_ucose** - These options are turned off by default. See Related Topics.
- **Inode to Pathname** - The Inode to Pathname function creates relationships between file names and relative paths. If Inode to Pathname is disabled on a volume, you must enable it, then delete existing snapshots on the volume. When new snapshots are created with Inode to Pathname enabled, the volume can be cataloged.

Internationalization Requirements

- The language code must be set and the UTF-8 variant must be specified on the NetApp storage system. For example, **en_US.UTF-8**. Only the English locale for vol0 for UTF-8 is supported. See Related Topics.
- The IBM Spectrum Copy Data Management application and documentation are available in English only. However, cataloging, searching, and reporting functions support international metadata.

Pure Storage Requirements

- FlashArray//c
- FlashArray//m
- FlashArray//x
- FlashArray-4xx series
- REST API version 1.16 and later
- Purity 5.2.X and later

VMware Requirements

- vSphere 6.0 and 6.0 update and patch levels
- vSphere 6.5 and 6.5 update and patch levels
- vSphere 6.7 and 6.7 update and patch levels

Ensure the latest version of VMware Tools is installed in your environment. IBM Spectrum Copy Data Management was tested against VMware Tools 9.10.0.

Third Party License Requirements

For NetApp ONTAP storage systems, a FlexClone license is required to fully utilize Instant Disk Restore and Instant VM Restore features. SnapVault and SnapMirror licenses are also required on both source and destination resources if your workflow includes vault and mirror protection. Similarly, a SnapRestore license is required if your workflow includes individual file recovery. Note that individual file recovery is only supported at the primary site.

For IBM storage systems, Remote Mirroring and FlashCopy licenses are required.

For VMware, a Storage vMotion license is required. A Kroll license is required for individual mailbox recovery on a Microsoft Exchange VM protected by IBM Spectrum Copy Data Management.

RELATED TOPICS:

- [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page **63**
- [File System Requirements](#) on page **28**
- [InterSystems Caché Requirements](#) on page **34**

- [Oracle Requirements](#) on page **41**
- [SAP HANA Requirements](#) on page **51**
- [Microsoft SQL Server Requirements](#) on page **57**
- [AWS Requirements](#) on page **62**
- [Oracle Database Support FAQ](#) on page **563**
- [Microsoft SQL Server Support FAQ](#) on page **584**
- [NetApp Support Article: Converting existing directories to Unicode format](#)
- [NetApp Knowledge Base Article 1010992: How to enable or disable I2P on a volume](#)
- [NetApp Knowledge Base Article 1013246: How to use UTF-8 characters with Solaris NFSv4](#)

File System Requirements

- [Supported Platforms/Configurations](#) on page 28
- [Software](#) on page 30
- [Connectivity](#) on page 30
- [Authentication](#) on page 31
- [Privileges](#) on page 31
- [Restore Jobs](#)
- [Revert Jobs](#) on page 32
- [Sample Configuration of an IBM Spectrum Copy Data Management Agent User](#) on page 32

Supported Platforms/Configurations

File System Configuration Requirements

Operating Systems	Server Types	Storage Systems	Storage Configuration
<ul style="list-style-type: none"> • Red Hat Enterprise Linux 6.5+ • Red Hat Enterprise Linux 7.0+ • SUSE Linux Enterprise Server 12.0+ • AIX 6.1+ • AIX 7.1+ • Windows Server 2012 R2 [1] • Windows Server 2016 [1] • Windows Server 2019 [1] 	Physical [4]	<ul style="list-style-type: none"> • HPE Nimble Storage 5.2 and later [6] [7] • IBM Spectrum Accelerate 11.5.3 and later: <ul style="list-style-type: none"> ◦ IBM FlashSystem A9000/A9000R [8] ◦ IBM XIV storage systems [8] • IBM Spectrum Virtualize Software 7.3 and later/8.1.2 and later: <ul style="list-style-type: none"> ◦ IBM SAN Volume Controller [8] ◦ IBM Storwize [8] ◦ IBM FlashSystem V9000 and 9100 [8] • Pure Storage running Pure APIs 1.5 and later: <ul style="list-style-type: none"> ◦ FlashArray//c ◦ FlashArray//m ◦ FlashArray//x ◦ FlashArray 4xx series 	<ul style="list-style-type: none"> • Fibre Channel [6] • iSCSI [6]
	Virtual (VMware) [2]	<ul style="list-style-type: none"> • DellEMC Unity 	<ul style="list-style-type: none"> • iSCSI disks

		<ul style="list-style-type: none"> ◦ EMC Unity 300, 400, 500, 600 (All-Flash and Hybrid Flash) ◦ EMC UnityVSA ◦ EMC VNXe 1600 running version 3.1.3 + ◦ EMC VNXe 3200 running version 3.1.1 + • IBM Spectrum Accelerate 11.5.3 and later: <ul style="list-style-type: none"> ◦ IBM FlashSystem A9000/A9000R [8] ◦ IBM XIV storage systems [8] • IBM Spectrum Virtualize Software 7.3 and later/8.1.2 and later: <ul style="list-style-type: none"> ◦ IBM SAN Volume Controller [8] ◦ IBM Storwize [8] ◦ IBM FlashSystem V9000 and 9100 [8] • NetApp ONTAP storage systems running the following versions: <ul style="list-style-type: none"> ◦ Data ONTAP 8.1.0, 8.2.0 and later in 7-mode ◦ Data ONTAP 9.x ◦ Clustered Data ONTAP 8.1, 8.2, 8.3, 9.4 and later • Pure Storage running Pure APIs 1.5 and later: <ul style="list-style-type: none"> ◦ FlashArray//c ◦ FlashArray//m ◦ FlashArray//x ◦ FlashArray 4xx series 	<ul style="list-style-type: none"> • attached directly to guest operating system [6] • pRDM [3]
--	--	--	---

[1] Windows Remote Shell (WinRM) must be enabled.

[2] See [System Requirements](#) on page 23 for supported VMware vSphere versions.

[3] Select the Physical provider type when registering the provider in IBM Spectrum Copy Data Management. Note that NetApp ONTAP and DellEMC storage systems are not supported. All data files and log files for a given instance should reside on either a pRDM or virtual disk.

[4] Dynamic disks are not supported.

[5] On AIX LPAR/VIO servers, data mount points must reside on disks attached to the server using NPIV. Virtual SCSI disks are not supported.

[6] HPE Nimble Storage must be version 5.2 or later to support iSCSI and Fibre Channel.

[7] The make permanent option is not available for HPE Nimble Storage using physical disks.

[8] On IBM Systems Storage, `condense` is run during maintenance jobs.

Windows File System Considerations

For Windows File System, the following configurations are not supported:

- Windows Failover Clusters
- Clustered Shared Volumes (CSV)

NTFS is required for those configurations.

Software

- The `bash` and `sudo` packages must be installed. `Sudo` must be version 1.7.6p2 or above. Run `sudo -V` to check the version.
- Python version 2.6.x or 2.7.x must be installed.
- **AIX only:** If data resides on IBM Spectrum Accelerate storage, the IBM Storage Host Attachment Kit (also known as IBM XIV Host Attachment Kit) must be installed on the server.
- **RHEL/OEL/CentOS 6.x only:** Ensure the `util-linux-ng` package is up-to-date by running `yum update util-linux-ng`. Depending on your version or distribution, the package may be named `util-linux`.
- **RHEL/OEL/CentOS 7.3 and above:** A required Perl module, `Digest::MD5`, is not installed by default. Install the module by running `yum install perl-Digest-MD5`.
- **Linux only:** If data resides on LVM volumes, ensure the LVM version is 2.0.2.118 or later. Run `lvm version` to check the version and run `yum update lvm2` to update the package if necessary.
- **Linux only:** If data resides on LVM volumes, the `lvm2-lvmetad` service must be disabled as it can interfere with IBM Spectrum Copy Data Management's ability to mount and resignature volume group snapshots/clones.

Run the following commands to stop and disable the service:

```
systemctl stop lvm2-lvmetad
systemctl disable lvm2-lvmetad
```

Additionally, disable `lvmetad` in the LVM config file. Edit the file `/etc/lvm/lvm.conf` and set:

```
use_lvmetad = 0
```

Connectivity

- The SSH service must be running on port 22 on the server and any firewalls must be configured to allow IBM Spectrum Copy Data Management to connect to the server using SSH. The SFTP subsystem for SSH must also be enabled.
- The server can be registered using a DNS name or IP address. DNS names must be resolvable by IBM Spectrum Copy Data Management.
- In order to mount clones/copies of data, IBM Spectrum Copy Data Management automatically maps and unmaps LUNs to the servers. Each server must be preconfigured to connect to the relevant storage systems at that site.
 - For Fibre Channel, the appropriate zoning must be configured beforehand.
 - For iSCSI, the servers must be configured beforehand to discover and log in to the targets on the storage servers.

Authentication

- The application server must be registered in IBM Spectrum Copy Data Management using an operating system user that exists on the server (referred to as "IBM Spectrum Copy Data Management agent user" for the rest of this topic).
- During registration you must provide either a password or a private SSH key that IBM Spectrum Copy Data Management will use to log in to the server.
- For password-based authentication ensure the password is correctly configured and that the user can log in without facing any other prompts, such as prompts to reset the password.
- For key-based authentication ensure the public SSH key is placed in the appropriate *authorized_keys* file for the IBM Spectrum Copy Data Management agent user.
 - Typically, the file is located at `/home/<username>/.ssh/authorized_keys`
 - Typically, the `.ssh` directory and all files under it must have their permissions set to 600.

Privileges

The IBM Spectrum Copy Data Management agent user must have the following privileges:

- Privileges to run commands as root and other users using `sudo`. IBM Spectrum Copy Data Management requires this for various tasks such as discovering storage layouts and mounting and unmounting disks.
 - The *sudoers* configuration must allow the IBM Spectrum Copy Data Management agent user to run commands without a password.
 - The *!requiretty* setting must be set.

For examples on creating a new user with the necessary privileges, see [Sample Configuration of an IBM Spectrum Copy Data Management Agent User](#) on page 32.

Restore Jobs

When performing a database or filesystem restore from an XFS filesystem, the restore process may fail if the "xfsprogs" package version on the destination server is between 3.2.0 and 4.1.9. To resolve the issue, upgrade xfsprogs to version 4.2.0 or above.

Revert Jobs

When performing File System revert jobs, there are considerations that must be met for jobs to successfully complete:

- Ensure that the databases are on independent storage. The underlying storage volume for the databases being reverted should not contain data for other databases and should not contain a datastore that is shared by other virtual machines (VMs) or by other databases not being reverted.
- Ensure that the production databases are not on VMDK disks that are part of a VMware VM snapshot.
- Please shut down any application databases on the volumes to be reverted prior to running revert.
- When the revert action completes, take appropriate steps to restore the applications to use the *old* data that is now contained on the volume(s) that were reverted.

Sample Configuration of an IBM Spectrum Copy Data Management Agent User

The commands below are examples for creating and configuring an operating system user that IBM Spectrum Copy Data Management will use to log in to the application server. The command syntax may vary depending on your operating system type and version.

- Create the user that will be designated as the IBM Spectrum Copy Data Management agent user: `useradd -m cdmagent`
- Set a password if using password-based authentication: `passwd cdmagent`
- If using key-based authentication, place the public key in `/home/cdmagent/.ssh/authorized_keys`, or the appropriate file depending on your `sshd` configuration, and ensure the correct ownership and permissions are set, such as:

```
chown -R cdmagent:cdmagent /home/cdmagent/.ssh
chmod 700 /home/cdmagent/.ssh
chmod 600 /home/cdmagent/.ssh/authorized_keys
```

- Place the following lines at the end of your `sudoers` configuration file, typically `/etc/sudoers`. If your existing `sudoers` file is configured to import configuration from another directory (for example, `/etc/sudoers.d`), you can also place the lines in a new file in that directory:

```
Defaults:cdmagent !requiretty
cdmagent ALL=(ALL) NOPASSWD:ALL
```

RELATED TOPICS:

- [Register a Provider](#) on page **75**
- [Identities Overview](#) on page **127**
- [System Requirements](#) on page **23**

InterSystems Caché Requirements

Before registering each InterSystems Caché server in IBM Spectrum Copy Data Management, ensure it meets the following requirements.

- [Supported Platforms/Configurations](#) on page **34**
- [InterSystems Caché Support for VMware Virtual Machines](#) on page **35**
- [Software](#) on page **36**
- [Connectivity](#) on page **36**
- [Authentication](#) on page **37**
- [Sample Configuration of an IBM Spectrum Copy Data Management Agent User](#) on page **38**
- [Backup Jobs](#) on page **38**
- [Restore Jobs](#) on page **38**
- [Performing a Caché Instant Database \(DB\) Restore](#) on page **39**

Supported Platforms/Configurations

InterSystems Caché Configuration Requirements

Database Versions/Types	Server Types	Operating Systems	Storage Configuration	Storage Systems
<ul style="list-style-type: none"> • Caché 2015+ [1] • Caché 2016+ [1] • Caché 2017+ [1] 	Physical	<ul style="list-style-type: none"> • AIX 6.1 TL9+ [5] • AIX 7.1+ [5] 	<ul style="list-style-type: none"> • Fibre Channel • iSCSI 	<ul style="list-style-type: none"> • IBM Spectrum Virtualize Software 7.3 and later/8.1.2 and later: <ul style="list-style-type: none"> ◦ IBM SAN Volume Controller [6] ◦ IBM Storwize [6] ◦ IBM FlashSystem V9000 and 9100 [6] • IBM Spectrum Accelerate 11.5.3 and later: <ul style="list-style-type: none"> ◦ IBM FlashSystem A9000/A9000R [6] ◦ IBM XIV storage systems [6]

	Virtual (VMware) [2, 4]	<ul style="list-style-type: none"> Red Hat Enterprise Linux 6.5+ 	<ul style="list-style-type: none"> Physical RDM backed by Fibre Channel or iSCSI disks attached to ESX [3] VMDK (dependent as well as independent disks) on VMFS datastores backed by Fibre Channel or iSCSI disks attached to ESX iSCSI disks directly attached to guest operating system [3] 	<ul style="list-style-type: none"> Pure Storage running Pure APIs 1.5 and later: <ul style="list-style-type: none"> FlashArray//c FlashArray//m FlashArray//x FlashArray 4xx series
--	-------------------------	---	---	---

[1] IBM Spectrum Copy Data Management was tested against Caché 2015.3 and 2017.1.

[2] See [System Requirements](#) on page 23 for supported VMware vSphere versions.

[3] Select the Physical provider type when registering the provider in IBM Spectrum Copy Data Management. Note that NetApp ONTAP and DellEMC storage systems are not supported. All data files and log files for a given instance should reside on either a pRDM or virtual disk.

[4] InterSystems Caché servers registered as virtual must have VMware Tools installed and running.

[5] Supported platforms: IBM Power Systems (Little Endian).

[6] On IBM Systems Storage, condense is run during maintenance jobs.

InterSystems Caché Support for VMware Virtual Machines

For virtual machines, the Disk UUID option must be enabled. To enable, power off the guest machine through the vSphere client, then select the guest and click Edit Settings. Select VM Options, then Edit Configuration under the Advanced section. Add a new row that states the following:

Name: `disk.enableUUID`

Value: `TRUE`

pRDM and virtual disks are supported for VMware configurations. Note that data files and log files for a given instance should reside on either pRDMs or virtual disks. One instance cannot utilize both disk types.

Cluster Support

For cluster configurations, register the application server in IBM Spectrum Copy Data Management using the cluster IP address. IBM Spectrum Copy Data Management will utilize the cluster IP during application discovery

and snapshot operations. In addition, the user configured in IBM Spectrum Copy Data Management with the cluster must be present on all cluster nodes and should preferably use the same UID and GID across all nodes.

Additionally, on each cluster node, create the file `/etc/guestapps.conf` with the following contents:

```
[unixagent]
overrideHostname = <CLUSTER NAME>
```

Software

- The *bash* and *sudo* packages must be installed. *Sudo* must be version 1.7.6p2 or above. Run `sudo -V` to check the version.
- Python version 2.6.x or 2.7.x must be installed.
- **AIX only:** If data resides on IBM Spectrum Accelerate storage, the IBM Storage Host Attachment Kit (also known as IBM XIV Host Attachment Kit) must be installed on the server.
- **RHEL/OEL/CentOS 6.x only:** Ensure the *util-linux-ng* package is up-to-date by running `yum update util-linux-ng`. Depending on your version or distribution, the package may be named *util-linux*.
- **RHEL/OEL/CentOS 7.3 and above:** A required Perl module, `Digest::MD5`, is not installed by default. Install the module by running `yum install perl-Digest-MD5`.
- **Linux only:** If data resides on LVM volumes, ensure the LVM version is 2.0.2.118 or later. Run `lvm version` to check the version and run `yum update lvm2` to update the package if necessary.
- **Linux only:** If data resides on LVM volumes, the *lvm2-lvmetad* service must be disabled as it can interfere with IBM Spectrum Copy Data Management's ability to mount and resignature volume group snapshots/clones.

Run the following commands to stop and disable the service:

```
systemctl stop lvm2-lvmetad
systemctl disable lvm2-lvmetad
```

Additionally, disable *lvmetad* in the LVM config file. Edit the file `/etc/lvm/lvm.conf` and set:

```
use_lvmetad = 0
```

Connectivity

- The SSH service must be running on port 22 on the server and any firewalls must be configured to allow IBM Spectrum Copy Data Management to connect to the server using SSH. The SFTP subsystem for SSH must also be enabled.
- The server can be registered using a DNS name or IP address. DNS names must be resolvable by IBM Spectrum Copy Data Management.

- In order to mount clones/copies of data, IBM Spectrum Copy Data Management automatically maps and unmaps LUNs to the servers. Each server must be preconfigured to connect to the relevant storage systems at that site.
 - For Fibre Channel, the appropriate zoning must be configured beforehand.
 - For iSCSI, the servers must be configured beforehand to discover and log in to the targets on the storage servers.

Authentication

- The application server must be registered in IBM Spectrum Copy Data Management using an operating system user that exists on the server (referred to as "IBM Spectrum Copy Data Management agent user" for the rest of this topic).
- During registration you must provide either a password or a private SSH key that IBM Spectrum Copy Data Management will use to log in to the server.
- For password-based authentication ensure the password is correctly configured and that the user can log in without facing any other prompts, such as prompts to reset the password.
- For key-based authentication ensure the public SSH key is placed in the appropriate *authorized_keys* file for the IBM Spectrum Copy Data Management agent user.
 - Typically, the file is located at `/home/<username>/.ssh/authorized_keys`
 - Typically, the `.ssh` directory and all files under it must have their permissions set to 600.

OS-level authentication must be enabled on the InterSystems Caché server. From the InterSystems Caché interface, navigate to:

System Administration > Security > System Security > Authentication/CSP Session Options, then select **Allow Operating System authentication**.

Next, navigate to **System Administration > Security > Services > %Service_Terminal**, then enable **Operating System** if the option is not already enabled.

The user identity associated with IBM Spectrum Copy Data Management registration must have sufficient privileges to invoke some system commands, such as those used to find disk information, as well as to run InterSystems Caché's "csession" command without asking for a database user name and password. To enable this feature, use the InterSystems Caché instance owner user after adding the user to "sudoers" file.

The IBM Spectrum Copy Data Management agent user must have the following privileges:

- Privileges to run commands as root and other users using sudo. IBM Spectrum Copy Data Management requires this for various tasks such as discovering storage layouts and mounting and unmounting disks.
 - The *sudoers* configuration must allow the IBM Spectrum Copy Data Management agent user to run commands without a password.
 - The *!requiretty* setting must be set.

For examples on creating a new user with the necessary privileges, see [Sample Configuration of an IBM Spectrum Copy Data Management Agent User](#) on page 38.

Sample Configuration of an IBM Spectrum Copy Data Management Agent User

The commands below are examples for creating and configuring an operating system user that IBM Spectrum Copy Data Management will use to log in to the application server. The command syntax may vary depending on your operating system type and version.

- Create the user that will be designated as the IBM Spectrum Copy Data Management agent user: `useradd -m cdmagent`
- Set a password if using password-based authentication: `passwd cdmagent`
- If using key-based authentication, place the public key in `/home/cdmagent/.ssh/authorized_keys`, or the appropriate file depending on your `sshd` configuration, and ensure the correct ownership and permissions are set, such as:

```
chown -R cdmagent:cdmagent /home/cdmagent/.ssh
chmod 700 /home/cdmagent/.ssh
chmod 600 /home/cdmagent/.ssh/authorized_keys
```

- Place the following lines at the end of your `sudoers` configuration file, typically `/etc/sudoers`. If your existing `sudoers` file is configured to import configuration from another directory (for example, `/etc/sudoers.d`), you can also place the lines in a new file in that directory:

```
Defaults:cdmagent !requiretty
cdmagent ALL=(ALL) NOPASSWD:ALL
```

Backup Jobs

InterSystems Caché backup jobs occur at the instance level. Multiple instances can be added to a single Backup job definition, and all instances on a given host are automatically discovered.

Note that it is possible to scan in an InterSystems Caché backup failover member instance or an async member instance and run snapshots against the mirror copy instead of the primary failover member.

Virtual instances can be selected for backup (such as "cachecentos6/cache_2015.3"), but it is recommended to select Caché instances explicitly. If virtual instances are selected, the associated job definition must be adjusted when the Caché instances are upgraded to later versions.

Restore Jobs

When performing a database or filesystem restore from an XFS filesystem, the restore process may fail if the "xfsprogs" package version on the destination server is between 3.2.0 and 4.1.9. To resolve the issue, upgrade xfsprogs to version 4.2.0 or above.

Caché software is not required on the target host. However, the target host should have similar specifications to the source host, including operating system and processor.

Note that the following users and groups must be created on the target host: instance owner, effective user for InterSystems Caché superserver and its jobs, effective group for InterSystems Caché processes, and a group that has permissions to start and stop InterSystems Caché instances. The user and group IDs should match those on the source host. The instance will be brought up using the same mount points as those found on the source machine, so ensure these mounts are not in use on the target.

When restoring to a target with running InterSystems Caché instances, the instances display as valid targets. Note that IBM Spectrum Copy Data Management will not interact with these instances, but instead bring up a new instance using mapped mount points. When restoring to a target with no prior InterSystems Caché instances, IBM Spectrum Copy Data Management creates a placeholder that acts as a restore target named *cache_general*. Note that *cache_general* should only be used as a restore target and should not be selected for backup.

Single InterSystems Caché databases can be restored through an Instant Disk Restore job, which mounts physical volumes on the target machine. Granular recovery can then be performed through InterSystems Caché commands.

Performing a Caché Instant Database (DB) Restore

Though not fully accessible through the IBM Spectrum Copy Data Management interface, you may need to perform a Caché instant database (DB) restore. This process requires access to the cache server system and use of the command-line interface (CLI).

Caché instant DB restore requires an instance directory `/cache` with `/mgr` and `/bin` as subdirectories. The `/bin` subdirectory must contain the *ccontrol* and *cstart* files. Finally, the instance directory must be copied from `/mnt/root` to `/`.

Follow the steps below to perform an instant database restore:

- Open a terminal on the cache server or use secure shell (SSH) to connect to the cache server. Default credentials for the cache server are `root` for the username and `<root password>` for the password.
- Create the instance directory `/cache` and create the two subdirectories `/bin` and `/mgr` within the instance directory.

```
$ mkdir /cache
```

```
$ cd cache
```

```
$ mkdir bin
```

```
$ mkdir mgr
```

- Copy *ccontrol* and *cstart* from any available different instance to the newly created `/cache/bin` location.

```
$ cp ccontrol /cache/bin
```

```
$ cp cstart /cache/bin
```

- Run the Caché instant restore job in IBM Spectrum Copy Data Management. Running the instant restore job will create the additional directories `/boot` and `/root` under the `/mnt` directory. Additionally, a `/cache` subdirectory will be created in the `/mnt/root` directory.
- Move the `/cache` directory created in the second step to another location and copy `/cache` from `/mnt/root` to the `/` location.

```
$ mv /cache /<some other location>
```

```
$ cp -r /mnt/root/cache /
```

- Bring up the restored Caché instance by issuing the following command:

```
$ cd /cache/bin
```

```
$ ccontrol start cache
```

RELATED TOPICS:

- [Register a Provider](#) on page 75
- [Identities Overview](#) on page 127
- [System Requirements](#) on page 23

Oracle Requirements

Before registering each Oracle server in IBM Spectrum Copy Data Management, ensure it meets the following requirements.

- [Oracle Configuration Requirements](#) on page **41**
- [Restore Jobs](#)
- [Software](#) on page **44**
- [Connectivity](#) on page **45**
- [Authentication](#) on page **45**
- [Privileges](#) on page **46**
- [Oracle Home Discovery for Symbolic Links](#) on page **47**
- [Database Discovery](#) on page **47**
- [ASM Disk Discovery](#) on page **48**
- [Sample Configuration of an IBM Spectrum Copy Data Management Agent User](#) on page **49**

Oracle Configuration Requirements

Oracle Configuration Requirements

Database Versions/Types	Server Types	Operating Systems	Storage Configuration	Storage Systems
Oracle 11g R2, 12c R2, or 18c [1] [13] configured as: <ul style="list-style-type: none"> • Standalone on ASM or File Systems [2, 3] • RAC on ASM [2, 3, 4, 5] Oracle 18c configured as: <ul style="list-style-type: none"> • Standalone or RAC on ACFS [2, 3, 4, 5, 16, 21] 	Physical [6]	<ul style="list-style-type: none"> • AIX 6.1 TL9+ [8] • AIX 7.1+ [8] • AIX 7.2+ [8, 21] 	<ul style="list-style-type: none"> • Fibre Channel [18] • iSCSI [18] • NFS 	<ul style="list-style-type: none"> • HPE Nimble Storage 5.2 and later [18, 19, and 20] • IBM Spectrum Virtualize Software 7.3 and later/8.1.2 and later: <ul style="list-style-type: none"> ◦ IBM SAN Volume Controller [22] ◦ IBM Storwize [22] ◦ IBM FlashSystem V9000 and 9100 [22] • IBM Spectrum Accelerate 11.5.3 and later:

<p>Oracle 19c configured as:</p>		<ul style="list-style-type: none"> • Red Hat Enterprise Linux / Centos / Oracle 6.5+ [9] • Red Hat Enterprise Linux / Centos / Oracle 7.0+ [9] • SUSE Linux Enterprise Server 11.0 SP4+ [9] • SUSE Linux Enterprise Server 12.0+ [9] 	<ul style="list-style-type: none"> • Fibre Channel [18] • iSCSI [18] • NFS 	<ul style="list-style-type: none"> ◦ IBM FlashSystem A9000/A9000R [22] ◦ IBM XIV storage systems [22] • NetApp ONTAP storage systems running the following versions [14][15]: <ul style="list-style-type: none"> ◦ Data ONTAP 8.1.0, 8.2.0 and later ◦ Data ONTAP 9.x ◦ Clustered Data ONTAP 8.1, 8.2, 8.3, 9.4, 9.5, 9.6 and later
<ul style="list-style-type: none"> • Standalone only [17] 	<p>Virtual (VMware) [6, 7, 10, 11, 12]</p>	<ul style="list-style-type: none"> • Red Hat Enterprise Linux / Centos / Oracle 6.5+ [9] • Red Hat Enterprise Linux / Centos / Oracle 7.0+ [9] • SUSE Linux Enterprise Server 11.0 SP4+ [9] • SUSE Linux Enterprise Server 12.0+ [9] 	<ul style="list-style-type: none"> • Physical RDM backed by Fibre Channel or iSCSI disks attached to ESX [18] • VMDK (dependent as well as independent disks) on VMFS datastores (Fibre Channel / iSCSI) or NFS datastores [18] • iSCSI disks directly attached to guest operating system [18] • NFS share mapped directly to the guest 	<ul style="list-style-type: none"> • Pure Storage running Pure APIs 1.5 and later: <ul style="list-style-type: none"> ◦ FlashArray//c ◦ FlashArray//m ◦ FlashArray//x ◦ FlashArray 4xx series

[1] For Oracle 12c multitenant databases, IBM Spectrum Copy Data Management supports protection and recovery of the container database, including all pluggable databases under it. Granular recovery of specific PDBs can be performed via Instant Disk Restore recovery combined with RMAN.

[2] Standalone databases protected by IBM Spectrum Copy Data Management can be recovered to the same or another standalone server as well as to a RAC installation. When recovering from standalone to RAC, if the source database uses Automatic Storage Management, then it will be successfully recovered to all nodes in the destination cluster. If the source database uses non-ASM storage, the database will be mounted only on the first node in the destination RAC. Source disks for standalone databases restoring to a virtual RAC environment must be thick provisioned.

RAC databases protected by IBM Spectrum Copy Data Management can be recovered to the same or another RAC installation as well as to a standalone server. In order to recover a RAC database to a standalone server, the destination server must have Grid Infrastructure installed and an ASM instance must be running.

[3] Oracle Data Guard primary and secondary databases support inventory and backup operations. Oracle Data Guard databases will always be restored as primary databases without any Data Guard configurations enabled.

[4] Oracle Flex ASM is not supported.

[5] RAC database recoveries are not server pool-aware. IBM Spectrum Copy Data Management can recover databases to a RAC, but not to specific server pools.

[6] IBM Spectrum Copy Data Management supports recovering databases from a source physical server to a destination virtual server by provisioning disks as physical RDMs. Similarly, IBM Spectrum Copy Data Management can recover databases from a source virtual server that uses physical RDM to a destination physical server. However, source databases on VMDK virtual disks can only be recovered to another virtual server and not to a physical server.

[7] IBM Spectrum Copy Data Management does not support VADP-based protection of virtual Oracle servers. Oracle data must reside directly on one of the supported storage systems listed above.

[8] On AIX LPAR/VIO servers, Oracle data must reside on disks attached to the server using NPIV. Virtual SCSI disks are not supported.

[9] Linux LVM volumes containing Oracle data must use LVM version 2.02.118 or above. On SLES 11 SP4, this version of LVM may not be available through the official repositories, in which case, databases running on or recovered to SLES 11 SP4 systems must use ASM or non-LVM filesystems only.

[10] Masking and DevOps recoveries are not supported on virtual servers.

[11] See [System Requirements](#) on page 23 for supported VMware vSphere versions.

[12] Oracle servers registered as virtual must have VMware Tools installed and running.

[13] Oracle 12c multithreaded configurations are not supported.

[14] Data masking is not supported for Oracle in NetApp storage environments. Masking is not supported on source database on NFS (a copy of which will be cloned and masked) or source databases on replica copies. Instead of the default mirror copy, you must select a snapshot copy as a replication source.

[15] NetApp systems running in 7-Mode are not supported.

[16] The `ORACLE_HOME` directory may reside on ACFS. The `data` and `log` directories are required to be on ACFS. IBM Spectrum Copy Data Management does not protect the Oracle installation that is contained in the `ORACLE_HOME` directory.

[17] Oracle 19c is only supported for standalone on AIX 7.2. It is not supported using RAC.

Additional Requirements:

- Oracle database data and the flash recovery area (FRA) must reside on supported storage systems. IBM Spectrum Copy Data Management can back up archived logs to a supported storage system if they are not already on one.

- During an Instant Database Restore, there may be failures if the new name specified for the restored database is similar to an existing database only differing by numerical suffix. For clustered instances of Oracle databases, the appliance always uses global database name in the UI. During the inventory and restore processes, individual instances using the numerical suffixes of the cluster must be correlated to the global database name. The issue with this comes when, as an example, "Production12" is discovered. Is this the instance 12 of the "Production" database, or instance two (2) of the "Production1" database, or a database named "Production12."

[18] HPE Nimble Storage must be version 5.2 or later to support iSCSI and Fibre Channel.

[19] Oracle PIT (point-in-time) recovery is not supported with HPE Nimble Storage.

[20] The make permanent option is not available for HPE Nimble Storage using physical disks.

[21] Oracle ACFS on AIX 7.2 is supported. Oracle ACFS on AIX is not supported for flashcopy incremental backups and iSCSI protocol.

[22] On IBM Systems Storage, condense is run during maintenance jobs.

For more information about Oracle requirements, see [Oracle Database Support FAQ](#) on page 563.

Oracle Support for VMware Virtual Machines

For Oracle servers running as VMware virtual machines, UUID must be enabled to perform Oracle-based Backup functions. To enable, power off the guest machine through the vSphere client, then select the guest and click **Edit Settings**. Select **Options**, then **General** under the Advanced section. Select **Configuration Parameters...**, then find the *disk.EnableUUID* parameter. If set to FALSE, change the value to TRUE. If the parameter is not available, add it by clicking **Add Row**, set the value to TRUE, then power on the guest.

Oracle support for VMware virtual machines requires Oracle data/logs to be stored on VMDK virtual disks or physical RDMs. Virtual RDM disks are not supported. The VMDKs must reside on a datastore created on LUNs from supported storage systems. Similarly, the physical RDMs must be backed by LUNs from supported storage systems.

For Oracle RAC clustered nodes running prior to vSphere 6.0, virtual machines cannot use a virtual SCSI controller whose SCSI Bus Sharing option is set to None. This is required to ensure that IBM Spectrum Copy Data Management can hot-add shared virtual disks to the cluster nodes. For vSphere 6.0 and above, this requirement does not apply. Instead, if an existing shared SCSI controller is not found on vSphere 6.0 and above, IBM Spectrum Copy Data Management automatically enables the "multi-writer" sharing option for each shared virtual disk.

Restore Jobs

When performing a database or filesystem restore from an XFS filesystem, the restore process may fail if the "xfsprogs" package version on the destination server is between 3.2.0 and 4.1.9. To resolve the issue, upgrade xfsprogs to version 4.2.0 or above.

Software

- The *bash* and *sudo* packages must be installed. *Sudo* must be version 1.7.6p2 or above. Run `sudo -V` to check the version.
- Python version 2.6.x or 2.7.x must be installed.
- **AIX only:** If Oracle data resides on IBM Spectrum Accelerate storage, the IBM Storage Host Attachment Kit (also known as IBM XIV Host Attachment Kit) must be installed on the Oracle server.
- **RHEL/OEL/CentOS 6.x only:** Ensure the *util-linux-ng* package is up-to-date by running `yum update util-linux-ng`. Depending on your version or distribution, the package may be named *util-linux*.
- **RHEL/OEL/CentOS 7.3 and above:** A required Perl module, `Digest::MD5`, is not installed by default. Install the module by running `yum install perl-Digest-MD5`.
- **Linux only:** If Oracle data resides on LVM volumes, ensure the LVM version is 2.0.2.118 or later. Run `lvm version` to check the version and run `yum update lvm2` to update the package if necessary.
- **Linux only:** If Oracle data resides on LVM volumes, the *lvm2-lvmetad* service must be disabled as it can interfere with IBM Spectrum Copy Data Management's ability to mount and resigature volume group snapshots/clones.

Run the following commands to stop and disable the service:

```
$ systemctl stop lvm2-lvmetad
$ systemctl disable lvm2-lvmetad
```

Additionally, disable *lvmetad* in the LVM config file. Edit the file `/etc/lvm/lvm.conf` and set:

```
use_lvmetad = 0
```

Connectivity

- The SSH service must be running on port 22 on the server and any firewalls must be configured to allow IBM Spectrum Copy Data Management to connect to the server using SSH. The SFTP subsystem for SSH must also be enabled.
- The server can be registered using a DNS name or IP address. DNS names must be resolvable by IBM Spectrum Copy Data Management.
- When registering Oracle RAC nodes, register each node using its physical IP or name. Do not use a virtual name or Single Client Access Name (SCAN).
- In order to mount clones/copies of Oracle data, IBM Spectrum Copy Data Management automatically maps and unmaps LUNs to the Oracle servers. Each server must be preconfigured to connect to the relevant storage systems at that site.
 - For Fibre Channel, the appropriate zoning must be configured beforehand.
 - For iSCSI, the Oracle servers must be configured beforehand to discover and log in to the targets on the storage servers.

Authentication

- The Oracle server must be registered in IBM Spectrum Copy Data Management using an operating system user that exists on the Oracle server (referred to as "IBM Spectrum Copy Data Management agent user" for the rest of this topic).
- During registration you must provide either a password or a private SSH key that IBM Spectrum Copy Data Management will use to log in to the server.
- For password-based authentication ensure the password is correctly configured and that the user can log in without facing any other prompts, such as prompts to reset the password.
- For key-based authentication ensure the public SSH key is placed in the appropriate *authorized_keys* file for the IBM Spectrum Copy Data Management agent user.
 - Typically, the file is located at `/home/<username>/.ssh/authorized_keys`
 - Typically, the `.ssh` directory and all files under it must have their permissions set to 600.

Privileges

The IBM Spectrum Copy Data Management agent user must have the following privileges:

- Privileges to run commands as root and other users using `sudo`. IBM Spectrum Copy Data Management requires this for various tasks such as discovering storage layouts and mounting and unmounting disks.
 - The *sudoers* configuration must allow the IBM Spectrum Copy Data Management agent user to run commands without a password.
 - The *!requiretty* setting must be set.
 - The *ENV_KEEP* setting must allow the `ORACLE_HOME` and `ORACLE_SID` environment variables to be retained.
- Privileges to read the Oracle inventory. IBM Spectrum Copy Data Management requires this to discover and collect information about Oracle homes and databases.
 - To achieve this, the IBM Spectrum Copy Data Management agent user must belong to the Oracle inventory group, typically named *oinstall*.
- SYSDBA privileges for database instances. IBM Spectrum Copy Data Management needs to perform database tasks like querying instance details, hot backup, RMAN cataloging, as well as starting/stopping instances during recovery.
 - To achieve this, the IBM Spectrum Copy Data Management agent user must belong to the OSDBA operating system group, typically named *dba*.
 - In the case of multiple Oracle homes each with a different OSDBA group, the IBM Spectrum Copy Data Management agent user must belong to each group.

- SYSASM privileges, if Automatic Storage Management (ASM) is installed. IBM Spectrum Copy Data Management needs to perform storage tasks like querying ASM disk information, as well as renaming, mounting, and unmounting diskgroups.
 - To achieve this, the IBM Spectrum Copy Data Management agent user must belong to the OSASM operating system group, typically named *asmadmin*.
- Shell user limits for the IBM Spectrum Copy Data Management agent user must be the same as those for the user that owns the Oracle home, typically named *oracle*. Refer to Oracle documentation for requirements and instructions on setting shell limits. Run `ulimit -a` as both the oracle user and the IBM Spectrum Copy Data Management agent user and ensure their settings are identical.

For examples on creating a new user with the necessary privileges, see [Sample Configuration of an IBM Spectrum Copy Data Management Agent User](#) on page 49.

Oracle Home Discovery for Symbolic Links

IBM Spectrum Copy Data Management discovers Oracle installations and databases by looking through the files `/etc/oraInst.loc` and `/etc/oratab`, as well as the list of running Oracle processes. When connecting to a database instance to discover its properties, IBM Spectrum Copy Data Management connects by setting the `ORACLE_HOME` environment variable based on the path that was auto-discovered. In some cases, when Oracle databases are used in combination with SAP software for instance, a symbolic link to the Oracle Home path may be created. This link path can then be used by the database administrator to connect to the database.

Consider the following example where the real Oracle Home path can be

`/u01/app/oracle/product/1210/dbhome_1`, while the symbolic link that points to it is `/oracle/PRODDB/1210`.

```
$ ls -l /oracle/PRODB
lrwxrwxrwx. 1 oracle oinstall 1210 -> /u01/app/oracle/product/1210/dbhome_1
```

When IBM Spectrum Copy Data Management connects to the instance using the real, auto-discovered path, queries against the database can fail due to the `ORACLE_HOME` environment variable being incorrectly set. To override the auto-discovered path and force IBM Spectrum Copy Data Management to use the correct symbolic link path, a configuration file must be defined on each Oracle server where symbolic links are in use.

Create or edit if it already exists a file named `/etc/guestapps_oraHomes.conf`. Insert into this file these entries:

```
# Define one entry per line.
# Each entry must be in the form: <instanceName> = <oraHomePath>
# Lines beginning with the '#' character are ignored.
PRODDB = /oracle/PRODDB/1210
TESTDB = /oracle/TESTDB/1210
```

To ensure that the configuration file is readable by the agent user, run the following command:

```
chmod 644 /etc/guestapps_oraHomes.conf
```

Database Discovery

IBM Spectrum Copy Data Management discovers Oracle installations and databases by looking through the files `/etc/oraInst.loc` and `/etc/oratab`, as well as the list of running Oracle processes. If the files are not present in their default location, the "locate" utility must be installed on the system so that IBM Spectrum Copy Data Management can search for alternate locations of these files.

IBM Spectrum Copy Data Management discovers databases and their storage layouts by connecting to running instances and querying the locations of their datafiles, log files, etc. In order for IBM Spectrum Copy Data Management to correctly discover databases during cataloging and copy operations, databases must be in "MOUNTED," "READ ONLY," or "READ WRITE" mode. IBM Spectrum Copy Data Management cannot discover or protect database instances that are shut down.

Databases must be started using a server parameter file (spfile). IBM Spectrum Copy Data Management does not support copy operations for databases that are started using a text-based parameter file (pfile).

ASM Disk Discovery

When IBM Spectrum Copy Data Management mounts snapshots/clones of ASM disks, it configures the disks to set the appropriate permissions required to make them discoverable by ASM:

- The disk owner and group are set to the owner of the Grid installation and the OSASM group respectively. These are typically *grid* and *asmadmin*. IBM Spectrum Copy Data Management automatically discovers the appropriate owner and group information on each server.
- The disk permissions are set to 660.

Additionally, IBM Spectrum Copy Data Management creates aliases/symbolic links with names that follow a consistent pattern. To ensure that ASM is able to discover the disks mapped by IBM Spectrum Copy Data Management, you must update the `ASM_DISKSTRING` parameter to add this pattern.

Linux:

IBM Spectrum Copy Data Management creates *udev* rules for each disk to set the appropriate ownership and permissions. The *udev* rules also create symbolic links of the form `/dev/ecx-asmdisk/<diskId>` that point to the appropriate device under `/dev`.

To ensure the disks are discoverable by ASM, add the following pattern to your existing `ASM_DISKSTRING`:
`/dev/ecx-asmdisk/*`

AIX:

IBM Spectrum Copy Data Management creates a device node (using `mknod`) of the form `/dev/ecx_asm<diskId>` that points to the appropriate `hdisk` under `/dev`. IBM Spectrum Copy Data Management also sets the appropriate ownership and permissions for this new device.

To ensure that the disks are discoverable by ASM, add the following pattern to your existing `ASM_DISKSTRING`: `/dev/ecx_asm*`

Notes:

- If the existing value of the `ASM_DISKSTRING` is empty, you may have to first set it to an appropriate value that matches all existing disks, then append the value above.

- If the existing value of the `ASM_DISKSTRING` is broad enough to discover all disks (for example, `/dev/*`), you may not need to update it.
- Refer to Oracle documentation for details about retrieving and modifying the `ASM_DISKSTRING` parameter.

Sample Configuration of an IBM Spectrum Copy Data Management Agent User

The commands below are examples for creating and configuring an operating system user that IBM Spectrum Copy Data Management will use to log in to the Oracle server. The command syntax may vary depending on your operating system type and version.

- Create the user that will be designated as the IBM Spectrum Copy Data Management agent user: `useradd -m cdmagent`
- Set a password if using password-based authentication: `passwd cdmagent`
- If using key-based authentication, place the public key in `/home/cdmagent/.ssh/authorized_keys`, or the appropriate file depending on your `sshd` configuration, and ensure the correct ownership and permissions are set, such as:


```
chown -R cdmagent:cdmagent /home/cdmagent/.ssh
chmod 700 /home/cdmagent/.ssh
chmod 600 /home/cdmagent/.ssh/authorized_keys
```
- Add the user to the Oracle installation and OSDBA group: `usermod -a -G oinstall,dba cdmagent`
- If ASM is in use, also add the user to the OSASM group: `usermod -a -G asmadmin cdmagent`

Note: If on AIX, the append argument (`-a`) should be omitted when using the `usermod` command.

- Place the following lines at the end of your `sudoers` configuration file, typically `/etc/sudoers`. If your existing `sudoers` file is configured to import configuration from another directory (for example, `/etc/sudoers.d`), you can also place the lines in a new file in that directory:

```
Defaults:cdmagent !requiretty
Defaults:cdmagent env_keep+="ORACLE_HOME"
Defaults:cdmagent env_keep+="ORACLE_SID"
cdmagent ALL=(ALL) NOPASSWD:ALL
```

RELATED TOPICS:

- [Oracle Database Support FAQ](#) on page 563
- [Register a Provider](#) on page 75

- [Identities Overview](#) on page **127**
- [System Requirements](#) on page **23**

SAP HANA Requirements

Before registering each SAP HANA server in IBM Spectrum Copy Data Management, ensure it meets the following requirements.

- [Supported Platforms/Configurations](#) on page **51**
- [Prerequisites](#) on page **52**
- [Software](#) on page **53**
- [Connectivity](#) on page **54**
- [Authentication](#) on page **54**
- [Registration](#) on page **54**
- [Privileges](#) on page **55**
- [Restore Jobs](#) on page **55**
- [Revert Jobs](#) on page **55**
- [Sample Configuration of an IBM Spectrum Copy Data Management Agent User](#) on page **56**

Supported Platforms/Configurations

SAP HANA Configuration Requirements

Database Versions/Types	Server Types	Operating Systems	Storage Configuration	Storage Systems
<ul style="list-style-type: none"> • SAP HANA 1.0 SPS 07+ [9] • SAP HANA 2.0 SPS 02 [5, 9] • SAP HANA 2.0 SPS 03 [9] • SAP HANA 2.0 SPS 04 [9] (See Prerequisite note below.) <p>Note that only one tenant is supported in Multiple Database</p>	Physical [4]	<ul style="list-style-type: none"> • Red Hat Enterprise Linux 6.5+ • Red Hat Enterprise Linux 7.0+ • SUSE Linux Enterprise Server 12.0+ 	<ul style="list-style-type: none"> • Fibre Channel [6] • iSCSI [6] 	<ul style="list-style-type: none"> • HPE Nimble Storage 5.2 and later [6, 7, and 8] • IBM Spectrum Virtualize Software 7.3 and later/8.1.2 and later: <ul style="list-style-type: none"> ◦ IBM SAN Volume Controller [10] ◦ IBM Storwize [10] ◦ IBM FlashSystem V9000 and 9100 [10] • IBM Spectrum Accelerate 11.5.3 and later:

<p>Container (MDC) setups, in addition to the system tenant.</p>	<p>Virtual (VMware) [1, 3, 4]</p>	<ul style="list-style-type: none"> • Red Hat Enterprise Linux 6.5+ • Red Hat Enterprise Linux 7.0+ • SUSE Linux Enterprise Server 12.0+ 	<ul style="list-style-type: none"> • Physical RDM backed by Fibre Channel or iSCSI disks attached to ESX [2, 6] • VMDK (dependent as well as independent disks) on VMFS datastores backed by Fibre Channel or iSCSI disks attached to ESX [6] • iSCSI disks directly attached to guest operating system [2, 6] 	<ul style="list-style-type: none"> ◦ IBM FlashSystem A9000/A9000R [10] ◦ IBM XIV storage systems [10] • Pure Storage running Pure APIs 1.5 and later: <ul style="list-style-type: none"> ◦ FlashArray//c ◦ FlashArray//m ◦ FlashArray//x ◦ FlashArray 4xx series
--	-----------------------------------	--	---	--

[1] See [System Requirements](#) on page 23 for supported VMware vSphere versions.

[2] Select the Physical provider type when registering the provider in IBM Spectrum Copy Data Management. Note that NetApp ONTAP and DellEMC storage systems are not supported.

[3] SAP HANA servers registered as virtual must have VMware Tools installed and running.

[4] Supported platforms: Intel based x86, IBM Power Systems (Little Endian)

[5] Single tenant configurations can be automatically protected using storage snapshots.

[6] HPE Nimble Storage must be version 5.2 or later to support iSCSI and Fibre Channel.

[7] SAP HANA revert is not supported with HPE Nimble Storage.

[8] The make permanent option is not available for HPE Nimble Storage using physical disks.

[9] Only the xfs filesystem is supported for SAP HANA data and logs locations.

[10] On IBM Systems Storage, condense is run during maintenance jobs.

Prerequisites

- The SAP HANA Client must be installed on your SAP HANA machine.
- Create a symbolic link to the SAP HANA Client installation directory through the following command:

```
ln -s <installation directory of SAP HANA Client> /opt/hana.
```

For example, if SAP HANA Client is installed in /hana/shared/<SID>/hdbclient, you would enter the following: `ln -s /hana/shared/<SID>/hdbclient/ /opt/hana.`

- For SAP Hana SPS 04, the `hdbcli` module must be installed. The `hdbcli` module should only be installed after the complete installation of the HANA client. The module may be extracted from `<path to directory>/hdbclient/hdbcli-<version>.tar.gz`.
- Each SAP HANA system has a system ID (SID). It is good practice to have the SID in the path. For example, if `/hana/logbackup` is the mount point, create these directories:

```
mkdir /hana/logbackup/<SID>
mkdir /hana/logbackup/<SID>/catalog
```

To ensure that database logs can be saved, permissions need to be appropriately specified for the created directories:

```
chown --reference=/hana/shared/<SID>/HDB01 /hana/logbackup/<SID>
chmod --reference=/hana/shared/<SID>/HDB01 /hana/logbackup/<SID>
chown --reference=/hana/shared/<SID>/HDB01 /hana/logbackup/<SID>/catalog
chmod --reference=/hana/shared/<SID>/HDB01 /hana/logbackup/<SID>/catalog
```

Software

- The `bash` and `sudo` packages must be installed. `Sudo` must be version 1.7.6p2 or above. Run `sudo -V` to check the version.
- Python version 2.6.x or 2.7.x must be installed.
- **RHEL/OEL/CentOS 6.x only:** Ensure the `util-linux-ng` package is up-to-date by running `yum update util-linux-ng`. Depending on your version or distribution, the package may be named `util-linux`.
- **RHEL/OEL/CentOS 7.3 and above:** A required Perl module, `Digest::MD5`, is not installed by default. Install the module by running `yum install perl-Digest-MD5`.
- **SLES only:** The Python `pip` package will need to be installed. Follow these steps to install the `pip` module on **SLES**:

```
$ sudo zypper addrepo
https://download.opensuse.org/repositories/Cloud:Tools/SLE_12_
SP4/Cloud:Tools.repo
$ sudo zypper refresh
$ sudo zypper python-pip
```

Note: The SLES equivalent of `yum` is `zypper`.

- **Linux only:** If data resides on LVM volumes, ensure the LVM version is 2.0.2.118 or later. Run `lvm version` to check the version and run `yum update lvm2` to update the package if necessary.
- **Linux only:** If data resides on LVM volumes, the `lvm2-lvmetad` service must be disabled as it can interfere with IBM Spectrum Copy Data Management's ability to mount and resignature volume group snapshots/clones.

Run the following commands to stop and disable the service:

```
systemctl stop lvm2-lvmetad
systemctl disable lvm2-lvmetad
```

Additionally, disable *lvmetad* in the LVM config file. Edit the file `/etc/lvm/lvm.conf` and set:

```
use_lvmetad = 0
```

Connectivity

- The SSH service must be running on port 22 on the server and any firewalls must be configured to allow IBM Spectrum Copy Data Management to connect to the server using SSH. The SFTP subsystem for SSH must also be enabled.
- The server can be registered using a DNS name or IP address. DNS names must be resolvable by IBM Spectrum Copy Data Management.
- In order to mount clones/copies of data, IBM Spectrum Copy Data Management automatically maps and unmaps LUNs to the servers. Each server must be preconfigured to connect to the relevant storage systems at that site.
 - For Fibre Channel, the appropriate zoning must be configured beforehand.
 - For iSCSI, the servers must be configured beforehand to discover and log in to the targets on the storage servers.

Authentication

- The application server must be registered in IBM Spectrum Copy Data Management using an operating system user that exists on the server (referred to as "IBM Spectrum Copy Data Management agent user" for the rest of this topic).
- During registration you must provide either a password or a private SSH key that IBM Spectrum Copy Data Management will use to log in to the server.
- For password-based authentication ensure the password is correctly configured and that the user can log in without facing any other prompts, such as prompts to reset the password.
- For key-based authentication ensure the public SSH key is placed in the appropriate *authorized_keys* file for the IBM Spectrum Copy Data Management agent user.
 - Typically, the file is located at `/home/<username>/.ssh/authorized_keys`
 - Typically, the `.ssh` directory and all files under it must have their permissions set to 600.

Registration

When registering a SAP HANA provider in IBM Spectrum Copy Data Management, note the following:

- The format for the port number is 3<instance number>15. So, for example, if the instance number is 07, then enter the following port number: 30715.
- The System Credentials are the credentials required to access the system. These credentials are used to log in to the system and perform restore operations.
- The Database Credentials are used for querying the database and performing copy operations. The user must have sufficient privileges to access the database.

Privileges

The IBM Spectrum Copy Data Management agent user must have the following privileges:

- Privileges to run commands as root and other users using sudo. IBM Spectrum Copy Data Management requires this for various tasks such as discovering storage layouts and mounting and unmounting disks.
 - The *sudoers* configuration must allow the IBM Spectrum Copy Data Management agent user to run commands without a password.
 - The *!requiretty* setting must be set.

For examples on creating a new user with the necessary privileges, see [SAP HANA Requirements](#) on page 51.

Restore Jobs

When performing a database or filesystem restore from an XFS filesystem, the restore process may fail if the "xfsprogs" package version on the destination server is between 3.2.0 and 4.1.9. To resolve the issue, upgrade xfsprogs to version 4.2.0 or above.

Revert Jobs

When performing revert jobs on SAP Hana, there are special considerations that must be met for jobs to successfully complete:

- SAP Hana data and the operating system (OS) filesystem must be on separate datastores.
- Ensure that the databases are on independent storage. The underlying storage volume for the databases being reverted should not contain data for other databases and should not contain a datastore that is shared by other virtual machines (VMs) or by other databases not being reverted.
- Ensure that the production databases are not on VMDK disks that are part of a VMware VM snapshot.
- All VM snapshots need to be removed from the SAP Hana server prior to executing the revert function.
- Production databases will be automatically shut down during the revert.
- Revert is available only once a restore has been completed.

When creating a job, you will need to set the default revert action for the job. Only SAP Hana can utilize the revert database functionality. This behavior is controlled through the Revert Database option during the job creation process:

- **Enabled** – Always will revert the database.
- **Disabled** – Never reverts the database.
- **User Selection** – Allows the user to make the determination to revert the database when the job session is pending.

Sample Configuration of an IBM Spectrum Copy Data Management Agent User

The commands below are examples for creating and configuring an operating system user that IBM Spectrum Copy Data Management will use to log in to the application server. The command syntax may vary depending on your operating system type and version.

- Create the user that will be designated as the IBM Spectrum Copy Data Management agent user: `useradd -m cdmagent`
- Set a password if using password-based authentication: `passwd cdmagent`
- If using key-based authentication, place the public key in `/home/cdmagent/.ssh/authorized_keys`, or the appropriate file depending on your `sshd` configuration, and ensure the correct ownership and permissions are set, such as:

```
chown -R cdmagent:cdmagent /home/cdmagent/.ssh
```

```
chmod 700 /home/cdmagent/.ssh
```

```
chmod 600 /home/cdmagent/.ssh/authorized_keys
```

- Place the following lines at the end of your `sudoers` configuration file, typically `/etc/sudoers`. If your existing `sudoers` file is configured to import configuration from another directory (for example, `/etc/sudoers.d`), you can also place the lines in a new file in that directory:

```
Defaults:cdmagent !requiretty
```

```
cdmagent ALL=(ALL) NOPASSWD:ALL
```

RELATED TOPICS:

- [Register a Provider](#) on page **75**
- [Identities Overview](#) on page **127**
- [System Requirements](#) on page **23**

Microsoft SQL Server Requirements

Before registering each SQL server in IBM Spectrum Copy Data Management, ensure it meets the following requirements.

- [Microsoft SQL Server Configuration Requirements](#) on page **57**
- [Registration and Authentication](#) on page **60**
- [Instant Seeding Prerequisites and Limitations](#) on page **60**
- [Privileges](#) on page **61**

Microsoft SQL Server Configuration Requirements

Microsoft SQL Server Configuration Requirements

Database Versions/Types and Operating Systems	Server Types	Storage Systems	Storage Configuration
<ul style="list-style-type: none"> • SQL 2012 on Windows Server 2012 R2 [1] • SQL 2012 SP2 on Windows Server 2016 • SQL 2014 on Windows Server 2012 R2 [1] • SQL 2014 on Windows Server 2016 • SQL 2016 on Windows Server 2012 R2 [1] • SQL 2016 on Windows Server 2016 • SQL 2017 on Windows Server 2019 <p>Standalone, SQL Server Failover Clustering, and AlwaysOn configurations are supported for</p>	<p>Physical [7] [8] [9] [10] [11]</p>	<ul style="list-style-type: none"> • HPE Nimble Storage 5.2 and later [12, 13, and 14] • IBM Spectrum Accelerate 11.5.3 and later: <ul style="list-style-type: none"> ◦ IBM FlashSystem A9000/A9000R [15] ◦ IBM XIV storage systems [15] • IBM Spectrum Virtualize Software 7.3 and later/8.1.2 and later: <ul style="list-style-type: none"> ◦ IBM SAN Volume Controller [15] ◦ IBM Storwize [15] ◦ IBM FlashSystem V9000 and 9100 [15] • Pure Storage running Pure APIs 1.5 and above: <ul style="list-style-type: none"> ◦ FlashArray//c ◦ FlashArray//m ◦ FlashArray//x ◦ FlashArray 4xx series 	<ul style="list-style-type: none"> • Fibre Channel [12] • iSCSI [12]

<p>SQL 2012, SQL 2014, and SQL 2016.[2]</p> <p>Note: It is highly recommended to install the latest SQL Server patches and updates in your environment.</p>	<p>Virtual (VMware) [3]</p>	<ul style="list-style-type: none"> • DellEMC Unity <ul style="list-style-type: none"> ◦ EMC Unity 300, 400, 500, 600 (All-Flash and Hybrid Flash) ◦ EMC UnityVSA ◦ EMC VNXe 1600 running version 3.1.3 + ◦ EMC VNXe 3200 running version 3.1.1 + • IBM Spectrum Accelerate 11.5.3 and later: <ul style="list-style-type: none"> ◦ IBM FlashSystem A9000/A9000R [15] ◦ IBM XIV storage systems [15] • IBM Spectrum Virtualize Software 7.3 and later/8.1.2 and later: <ul style="list-style-type: none"> ◦ IBM SAN Volume Controller [15] ◦ IBM Storwize [15] ◦ IBM FlashSystem V9000 and 9100 [15] • NetApp ONTAP storage systems running the following versions: <ul style="list-style-type: none"> ◦ Data ONTAP 8.1.0, 8.2.0 and later in 7-mode ◦ Data ONTAP 9.x ◦ Clustered Data ONTAP 8.1, 8.2, 8.3, 9.4 and later • Pure Storage running Pure APIs 1.5 and above: <ul style="list-style-type: none"> ◦ FlashArray//c ◦ FlashArray//m ◦ FlashArray//x ◦ FlashArray 4xx series <p>Note: VM Replication restore jobs can be run to store off-host copies on the storage systems listed above.</p>	<ul style="list-style-type: none"> • iSCSI disks attached directly to guest operating system [4, 12] • pRDM [4] • vRDM [5] • VMDK [6]
--	-----------------------------	---	---

[1] Windows Remote Shell (WinRM) must be enabled.

[2] Note that Clustered Shared Volumes (CSV) are not supported.

[3] See System Requirements for supported VMware vSphere versions.

[4] Select the Physical provider type when registering the provider in IBM Spectrum Copy Data Management. Recoveries require direct access to storage. Note that NetApp ONTAP and DellEMC storage systems are not supported.

[5] vRDMs are supported through VM Replication jobs.

[6] Independent disks are supported only if the underlying storage utilizes supported storage systems. Register the SQL resource as Physical when configuring the provider in IBM Spectrum Copy Data Management. Note that independent disks do not allow snapshots to be taken in VMware virtual scenarios. The above listed IBM Spectrum Accelerate, IBM Spectrum Virtualize, and Pure Storage FlashArrays are supported for physical registration.

[7] When registering physical SQL servers it is recommended to register via the DNS server. The IBM Spectrum Copy Data Management appliance must be resolvable and route-able by the DNS server; the physical SQL server will communicate back to IBM Spectrum Copy Data Management through DNS.

[8] Recovery for target servers registered as Physical provider types requires direct access to storage.

[9] Any Windows node with iSCSI or Fibre Channel access to the storage can be selected as a proxy server, provided that the node is not part of the original cluster. It is recommended to select a standalone virtual or physical Windows node as a proxy server.

[10] For physical SQL servers you must allow outgoing connections to port 8443 on the IBM Spectrum Copy Data Management appliance from the SQL server.

[11] Dynamic disks are not supported.

[12] HPE Nimble Storage must be version 5.2 or later to support iSCSI and Fibre Channel.

[13] The make permanent option is not available for HPE Nimble Storage using physical disks.

[14] SQL PIT (point-in-time) recovery is not supported with HPE Nimble Storage.

SQL servers residing on any storage can also be protected to supported storage systems through VM Replication jobs.

For both physical and virtual SQL environments, point-in-time recoveries beyond the last snapshot taken are incompatible with workflows utilizing more than one Site. In a virtual environment, the SQL server, associated vCenter, and storage must be registered to the same site. In a physical environment, the SQL server and storage must be registered to the same site.

[15] On IBM Systems Storage, condense is run during maintenance jobs.

For more information about Microsoft SQL Server requirements, see [Microsoft SQL Server Support FAQ](#) on page 584.

SQL Support for VMware Virtual Machines

UUID must be enabled to perform Microsoft SQL-based backup functions. To enable, power off the guest machine through the vSphere client, then select the guest and click **Edit Settings**. Select **Options**, then **General** under the Advanced section. Select **Configuration Parameters...**, then find the *disk.EnableUUID* parameter. If set to FALSE, change the value to TRUE. If the parameter is not available, add it by clicking **Add Row**, set the value to TRUE, then power on the guest.

The virtual machine must use SCSI disks only, dynamic disks are not supported.

The latest VMware Tools must be installed on the virtual machine node.

In-Memory OLTP Requirements and Limitations

In-Memory OLTP is a memory-optimized database engine used to improve database application performance, supported in SQL 2014 and 2016. Note the following IBM Spectrum Copy Data Management requirements and limitations for In-Memory OLTP usage:

- The maximum restore file path must be less than 256 characters, which is a SQL requirement. If the original path exceeds this length, consider using a customized restore file path to reduce the length.
- The metadata that can be restored is subject to VSS and SQL restore capabilities.

Instant Seeding Prerequisites and Limitations

The goal of seeding is to restore a secondary database by taking advantage of snapshot technology, and minimize the data movement between primary and secondary replicas.

Seeding of both virtual and physical SQL servers is supported. The seeding destination replica must be an instance of a secondary role and must be created or configured with a working mirroring endpoint before the seeding process. The seeding process restores using the original database only, and only the latest backup snapshot is supported.

Source databases must be backed up under the full recovery model, which is configured through the SQL Management Console. The full recovery model provides a granular transaction restore. Note that the transaction log can grow indefinitely unless the logs are backed up, which will automatically initiate log truncation. It is recommended to have a database administrator run maintenance to free space if necessary.

Before the seeding process, primary databases and log backups are required to ensure the LSN gap between the primary and secondary databases is acceptable by the SQL AAG framework.

The datafile and logfile paths on all replicas of an availability group must be the same. IBM Spectrum Copy Data Management requires that the original volume mount points are available on the target SQL node for seeding process. If the original volume was mounted on a volume mount point, such as a folder, the root drive letter of the mount point must exist before the restore begins.

When creating an Instant Seeding restore job definition, the destination must be a non-system drive. The SQL database and log files must be located on non-system drives.

A network share for Always On log backups must be accessible from the secondary node. The seeding restore will restore log backups directly from the original log backup target instead of the temporary snapshot.

Registration and Authentication

Register each SQL server as a provider in IBM Spectrum Copy Data Management by name or IP address. When registering a SQL Cluster (AlwaysOn) node, register each node by name or IP address. Note that the IP addresses must be public-facing and listening on port 5985. The fully qualified domain name and virtual machine node DNS name must be resolvable and route-able from the IBM Spectrum Copy Data Management appliance.

The user identity must have sufficient rights to install and start the IBM Spectrum Copy Data Management Tools Service on the node. This includes "Log on as a service" rights. For more information about the "Log on as a service" right, see <https://technet.microsoft.com/en-us/library/cc794944.aspx>.

The default security policy uses the Windows NTLM protocol, and the user identity format follows the default domain\Name format.

You must manually create a directory to store VSS provider logs when running IBM Spectrum Copy Data Management 2.2.6 and earlier. Create the following directory structure on the SQL server: c:\temp\CDM\logs

Kerberos Requirements

Kerberos-based authentication can be enabled through a configuration file on the IBM Spectrum Copy Data Management appliance. This will override the default Windows NTLM protocol.

For Kerberos-based authentication only, the user identity must be specified in the username@FQDN format. The username must be able to authenticate using the registered password to obtain a ticket-granting ticket (TGT) from the key distribution center (KDC) on the domain specified by the fully qualified domain name.

Kerberos authentication also requires that the clock skew between the Domain Controller and the IBM Spectrum Copy Data Management appliance is less than 5 minutes. Note that the default Windows NTLM protocol is not time dependent.

Privileges

On the SQL server, the system login credential must have **public** and **sysadmin** permissions enabled, plus permission to access cluster resources in a SQL AlwaysOn environment. If one user account is used for all SQL functions, a Windows login must be enabled for the SQL server, with **public** and **sysadmin** permissions enabled.

Every SQL instance can use a specific user account to access the resources of that particular SQL instance.

To perform log backups, the SQL user registered with IBM Spectrum Copy Data Management must have the **sysadmin** permission enabled to manage SQL server agent jobs. If the SQL server agent service user is the default NT user, the agent will use that account to enable/access log backup jobs.

RELATED TOPICS:

- [Microsoft SQL Server Support FAQ](#) on page **584**
- [Register a Provider](#) on page **75**
- [Identities Overview](#) on page **127**
- [System Requirements](#) on page **23**

AWS Requirements

Review the following requirements and pre-requisites for configuring an AWS provider for use with IBM Spectrum Copy Data Management.

IBM Spectrum Copy Data Management supports the on-premise cached storage gateway. An OVA is downloaded, deployed and activated through a wizard.

See docs.aws.amazon.com/storagegateway/latest/userguide/on-premises-gateway-common.html for more information.

For cache and upload buffer settings, best practices are defined in the following Amazon AWS topics:

docs.aws.amazon.com/storagegateway/latest/userguide/managing-cache-common.html

docs.aws.amazon.com/storagegateway/latest/userguide/managing-upload-buffer-common.html

Once the storage gateway is created with AWS and vCenter, register the AWS provider in IBM Spectrum Copy Data Management. IBM Spectrum Copy Data Management will automatically discover the storage gateway when a job is created. Note that 32 volumes are supported per storage gateway.

RELATED TOPICS:

- [Register a Provider](#) on page **75**
- [Identities Overview](#) on page **127**
- [System Requirements](#) on page **23**

Install IBM Spectrum Copy Data Management as a Virtual Appliance

To install IBM Spectrum Copy Data Management, deploy an OVF template. This creates a virtual appliance containing the application on a VMware host such as an ESX or ESXi server. To run IBM Spectrum Copy Data Management, access the newly created virtual machine.

BEFORE YOU BEGIN:

- Before deployment, run MD5 Checksum on the downloaded OVA file. Ensure the generated checksum matches the one provided in the MD5 Checksum file, which is part of the software download.
- You may need to configure an IP address pool associated with the VM network where you plan to deploy IBM Spectrum Copy Data Management. Correct configuration of the IP address pool includes the setup of IP address range (if used), netmask, gateway, DNS search string, and a DNS server IP address.
- To use DHCP instead of a static IP address, the network must have access to a properly configured DHCP server. Leave fields blank when prompted to enter network properties.
- To change the IP address allocation type after IBM Spectrum Copy Data Management deploys, redeploy the virtual machine.
- For later versions of vSphere, the vSphere Web Client may be required to deploy IBM Spectrum Copy Data Management appliances.
- Review the email from Data Protection Technical Support, which includes a product key necessary for running IBM Spectrum Copy Data Management.

To install IBM Spectrum Copy Data Management as a virtual appliance:

1. Use the vSphere Client to deploy IBM Spectrum Copy Data Management. From the **File** menu, choose **Deploy OVF Template**. If using the vSphere Web Client, click **Create/Register VM**, then select **Deploy a virtual machine from an OVF or OVA file**. Click **Next**.
2. Specify the location of the IBM Spectrum Copy Data Management OVA template file and select it. Click **Next**.
3. Review the template details and accept the End User License Agreement. Click **Next**.
4. Provide a meaningful name for the template, which becomes the name of your virtual machine. Identify an appropriate location to deploy the virtual machine. Click **Next**.
5. Identify the datacenter, server, and resource pool for deployment. When prompted to select storage, select from datastores already configured on the destination host. The virtual machine configuration file and virtual

disk files are stored on the datastore. Select a datastore large enough to accommodate the virtual machine and all of its virtual disk files. Click **Next**.

6. Select a disk format to store the virtual disks. It is recommended that you select thick provisioning, which is preselected for optimized performance. Thin provisioning requires less disk space, but may impact performance. Click **Next**.
7. Select networks for the deployed template to use. Several available networks on the ESX server may be available by clicking Destination Networks. Select a destination network that allows you to define the appropriate IP address allocation for the virtual machine deployment. Click **Next**.
8. Enter network properties for the virtual machine's default gateway, DNS, IP address and netmask. Leave fields blank to retrieve settings from a DHCP server. The virtual machine needs access to a DHCP server available on the configured destination network. Click **Next**.
9. Review your template selections. Click **Finish** to exit the wizard and to start deployment of the OVF template. Deployment might take significant time.
10. After OVF template deployment completes, power on your newly created virtual machine. This can be done from vSphere Client.

Note: The virtual machine must remain powered on for the IBM Spectrum Copy Data Management application to be accessible.

11. Make a note of the IP address of the newly created virtual machine. This is needed to log on to the application. Find the IP address in vSphere Client by clicking your newly created virtual machine and looking in the **Summary** tab.

Note: You must allow several minutes for IBM Spectrum Copy Data Management to initialize completely.

NEXT STEPS:

- Set your local time zone. See [Set Time Zone](#) on page 516.
- After the first use, you can enable additional users to logon by linking to an LDAP server.
- Start IBM Spectrum Copy Data Management and begin using it from any supported web browser. See [Start IBM Spectrum Copy Data Management](#) on page 65.

RELATED TOPICS:

- [Start IBM Spectrum Copy Data Management](#) on page 65

Start IBM Spectrum Copy Data Management

Launch IBM Spectrum Copy Data Management to begin using the application and its features.

BEFORE YOU BEGIN:

- IBM Spectrum Copy Data Management must be installed prior to starting the application. See [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page **63**.
- The System Administrator must provide you with the IP address for the virtual appliance and the IBM Spectrum Copy Data Management user name and password.

To start IBM Spectrum Copy Data Management:

1. From a supported browser, enter the following URL:

https://<HOSTNAME>:8443/portal/

where <HOSTNAME> is the IP address of the virtual machine where the application is deployed. This connects you to IBM Spectrum Copy Data Management.

2. In the logon dialog, enter your user name and password. If this is your first time logging on to IBM Spectrum Copy Data Management as a Super User, the default user name is **admin** and the default password is **password**. You will be prompted to reset the default Super User password.
3. Click **Sign In**. The application launches. The **Home** 🏠 tab displays the dashboard.

Note: You are automatically logged out of IBM Spectrum Copy Data Management after 30 minutes of inactivity. Log back in with your user name and password to continue.

NEXT STEPS:

- After the first use, enable additional users to logon by adding native users or linking to an LDAP server. See [Role-Based Access Control Overview](#) on page **104**.
- Add storage systems and virtual machine resources to the IBM Spectrum Copy Data Management database. See [Register a Provider](#) on page **75** and [Jobs Overview](#) on page **161**.
- Search or browse for objects that match certain criteria. See [Search for Objects](#) on page **346** and [Browse Inventory](#) on page **357**.
- Generate reports with predefined or customized parameters. See [Report Overview](#) on page **360**.

RELATED TOPICS:

- [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page **63**
- [System Requirements](#) on page **23**

Dashboard

The dashboard displays an overview of your IBM Spectrum Copy Data Management environment. Quickly review the status of your jobs and recently run reports.

BEFORE YOU BEGIN:

- Review IBM Spectrum Copy Data Management features and data flow information. See [IBM Spectrum Copy Data Management Overview](#) on page 10.

To re-enable widgets after closing and arrange them in their default position, click **Show All Widgets** . To collapse all widgets, click **Collapse All Widgets** . To expand all widgets, click **Expand All Widgets** .

Available Dashboard Widgets

Inventory Statistics

Displays an overview of the number of objects and the size of the data in each Catalog. Select a provider from the tabs to view the Inventory statistics for that object type. Hover over a bar in the graph to view its numeric value.

Appliance Filesystems

Displays an overview of the disk space used by the IBM Spectrum Copy Data Management appliances and provides appropriate warnings when space availability reaches a set threshold.

If the threshold is reached, identify jobs that are no longer in use and delete them, then run the Maintenance job to clean up resources. Alternatively, increase the disk space of your IBM Spectrum Copy Data Management data disks.

My Reports

Displays a list of generated reports including **Report Name**, **Date Generated**, and **Formats**.

My Jobs

Displays a list of defined jobs including **Job Name**, **Catalog Type**, and **Status**. Currently running job sessions are represented by a running icon. Once a job session finishes, one of the following icons appears in the status column:

-  **Completed** - Indicates the job session completed successfully. All tasks associated with the job session were completed.
-  **Partial** - Indicates the job session completed, but one or more tasks failed or were skipped.
-  **Failed** - Indicates the job session did not successfully complete due to mixed task statuses.

- ❌ **Aborted** - Indicates the job session did not successfully complete due to a reset, reboot, or shutdown of the virtual appliance server.
- ⏸ **Held** - Indicates the job has been paused through the Halt feature in the Actions menu.
- ⌚ **Idle** - Indicates the job session is idle.
- ⚠ **Skipped** - Indicates that a volume was not cataloged. See the Task tab for more information about skipped jobs.
- ⏹ **Stopped** - Indicates the job was stopped using the Stop button.

Job Success Rate

Displays a graph detailing the percentage of jobs that successfully completed in the past 10 days. Hover over a point in the line graph to view its numeric value.

NEXT STEPS:

- Review report details from the My Reports widget, such as available parameters and field definitions. Reports can be downloaded as HTML files, Adobe PDFs, Microsoft Excel spreadsheets, and Microsoft Word files.

RELATED TOPICS:

- [Sites & Providers Overview](#) on page **70**
- [Role-Based Access Control Overview](#) on page **104**
- [Jobs Overview](#) on page **161**
- [Monitor a Job Session](#) on page **164**
- [Search Overview](#) on page **345**
- [Report Overview](#) on page **360**

Sites & Providers

The topics in the following section cover adding, editing, and deleting sites and providers.

Sites & Providers Overview

Add Sites and providers IBM Spectrum Copy Data Management through the **Sites & Providers**  view on the **Configure**  tab.

- [Register a Provider](#) on page 75
- [Add a Site](#) on page 72

A site is a user-defined grouping of providers that is generally based on location, to help quickly identify and interact with data created through Copy Data Management jobs.

Sites are assigned when registering providers. When creating Backup and Restore jobs, sites clearly identify where your data is replicated by location.

Providers are physical servers that host objects and attributes. Once a provider is registered in IBM Spectrum Copy Data Management, cataloging, searching, and reporting can be performed.

Supported provider types are:

- Amazon Web Services (AWS). Supported types include Amazon Simple Storage Service (S3) cloud storage.
- Application servers. Supported application database types include InterSystems Caché, Oracle, SAP HANA, and SQL. Use the File System application type to register file systems for physical servers running Windows, Linux, and AIX. See associated application requirements for supported storage types.
- DellEMC storage systems. Supported types include DellEMC Unity.
- IBM storage systems. Supported types include IBM Spectrum Accelerate, IBM Spectrum Protect Snapshot, and IBM Spectrum Virtualize.
- LDAP servers. Register an LDAP server to enable LDAP users to be provisioned through a group import. LDAP also supports authentication using the sAMAccountName Windows user naming attribute or an associated e-mail address. Unlike storage systems and VMware servers, LDAP servers are not cataloged.
- NetApp ONTAP storage systems. Supported types include NetApp ONTAP 7-Mode and Cluster-Mode.
- SQL and Oracle application servers. Supported storage platforms for Oracle include IBM storage systems running IBM Spectrum™ Virtualize Software version 7.3 and later/8.1.2 and later, including IBM SAN Volume Controller, IBM Storwize, and IBM FlashSystem V9000 and 9100 systems. Supported SQL Server versions include SQL Server 2012, SQL Server 2014, SQL Server 2016 standalone and AlwaysON. Supported operating system platforms include Windows 2012R2, Windows 2016 running on vSphere VM using VMDK configuration.
- HPE Nimble Storage systems.
- Pure Storage systems. Supported storage platforms include Pure Storage FlashArray.

- SMTP hosts. Register an SMTP server to enable email notifications from IBM Spectrum Copy Data Management. Unlike storage systems and VMware servers, SMTP servers are not cataloged.
- VMware servers. Supported types includes vCenter and ESX/ESXi hosts.

Adding a provider requires specifying the user name and password of the provider.

Note: Users that register providers, such as storage devices, or add resources to IBM Spectrum Copy Data Management, such as jobs or customized reports, will have full access to interact with those providers or resources regardless of role-based access control restrictions. For example, if a user's permission allows them to register NetApp providers, they will also be able to view, edit, and unregister the NetApp providers that they registered, even if the necessary permissions are not assigned to them through role-based access control.

RELATED TOPICS:

- [Register a Provider](#) on page **75**
- [View a Provider](#) on page **95**
- [Edit a Provider](#) on page **98**
- [Unregister a Provider](#) on page **100**

Add a Site

Create sites to define a grouping of providers based on their location in your IBM Spectrum Copy Data Management environment. Once sites are created in IBM Spectrum Copy Data Management, they can be applied to your providers.

Click **View Relationship** to view the resources that are assigned to the site.

BEFORE YOU BEGIN:

- Review the properties and location of your current providers. See [Sites & Providers Overview](#) on page 70 for a list of supported providers.

To add a site:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Sites** tab.
The Sites pane opens.
2. In the Sites pane, click **New** . The Create Site dialog opens.
3. Enter a site name and a meaningful description.
4. To set this site as the default site, select **Set as default**. New storage providers are automatically assigned to the default site unless another site is selected.
5. Click **OK**. The site appears on the Sites pane and can be applied to new and existing storage providers.

NEXT STEPS:

- Assign sites to new and existing providers. See [Register a Provider](#) on page 75 and [Edit a Provider](#) on page 98.

RELATED TOPICS:

- [Edit a Site](#) on page 73
- [Delete a Site](#) on page 74

Edit a Site

Revise site names and descriptions to reflect location changes in your IBM Spectrum Copy Data Management environment.

To edit a site:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Sites** tab.
The Sites pane opens.
2. In the Sites pane, select the site to edit by clicking in the row containing the site name.
3. Click **Edit** . The Edit Site dialog opens.
4. Revise the site name and description.
5. To set this site as the default site, select **Set as default**. New storage providers are automatically assigned to the default site unless another site is selected.
6. Click **OK**. The revisions are applied to the site.

NEXT STEPS:

- Assign sites to new and existing providers. See [Register a Provider](#) on page 75 and [Edit a Provider](#) on page 98.

RELATED TOPICS:

- [Add a Site](#) on page 72
- [Delete a Site](#) on page 74

Delete a Site

Delete a site when it becomes obsolete.

A site cannot be deleted if it is assigned to a provider. On the Sites pane, click **View Relationship** to view the providers that are assigned to the site. Re-assign your providers to different sites before deleting.

To delete a site:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Sites** tab.
The Sites pane opens.
2. In the Sites pane, select the site to delete by clicking in the row containing the site name.
3. Click **Delete** . A confirmation dialog box displays.
4. Confirm deletion. The site is deleted.

NEXT STEPS:

- Assign sites to new and existing providers. See [Register a Provider](#) on page 75 and [Edit a Provider](#) on page 98.

RELATED TOPICS:

- [Add a Site](#) on page 72
- [Edit a Site](#) on page 73

Register a Provider

Providers are physical servers that host objects and attributes. Once a provider is registered in IBM Spectrum Copy Data Management, cataloging, searching, and reporting can be performed.

Supported provider types are:

- Amazon Web Services (AWS). Supported types include Amazon Simple Storage Service (S3) cloud storage.
- Application servers. Supported application database types include InterSystems Caché, Oracle, SAP HANA, and SQL. Use the File System application type to register file systems for physical servers running Windows, Linux, and AIX. See associated application requirements for supported storage types.
- DellEMC storage systems. Supported types include DellEMC Unity.
- IBM storage systems. Supported types include IBM Spectrum Accelerate, IBM Spectrum Protect Snapshot, and IBM Spectrum Virtualize.
- LDAP servers. Register an LDAP server to enable LDAP users to be provisioned through a group import. LDAP also supports authentication using the sAMAccountName Windows user naming attribute or an associated e-mail address. Unlike storage systems and VMware servers, LDAP servers are not cataloged.
- NetApp ONTAP storage systems. Supported types include NetApp ONTAP 7-Mode and Cluster-Mode.
- SQL and Oracle application servers. Supported storage platforms for Oracle include IBM storage systems running IBM Spectrum™ Virtualize Software version 7.3 and later/8.1.2 and later, including IBM SAN Volume Controller, IBM Storwize, and IBM FlashSystem V9000 and 9100 systems. Supported SQL Server versions include SQL Server 2012, SQL Server 2014, SQL Server 2016 standalone and AlwaysON. Supported operating system platforms include Windows 2012R2, Windows 2016 running on vSphere VM using VMDK configuration.
- HPE Nimble Storage systems.
- Pure Storage systems. Supported storage platforms include Pure Storage FlashArray.
- SMTP hosts. Register an SMTP server to enable email notifications from IBM Spectrum Copy Data Management. Unlike storage systems and VMware servers, SMTP servers are not cataloged.
- VMware servers. Supported types includes vCenter and ESX/ESXi hosts.

Adding a provider requires specifying the user name and password of the provider.

Storage providers can be automatically cataloged after registration. If the **Run Inventory job after registration** option is selected, IBM Spectrum Copy Data Management creates a high-level Inventory job and automatically catalogs the objects on the provider.

Note: Ensure that TLS protocol is enabled on the NetApp storage system by setting the `tls.enable` option to ON. For TLS to take effect on HTTPS, ensure that the `httpd.admin.ssl.enable` option is also set to ON. See [Enabling or disabling TLS](#) on NetApp's Support site.

BEFORE YOU BEGIN:

- Create a site to assign to your provider. A site is a user-defined grouping of providers that is generally based on location. See [Add a Site](#) on page 72.
- When registering providers it is recommended to assign all related sources, such as hosting vCenters and related storage systems, to the same Site.
- For application server it is recommended to keep all resources, such as hosting vCenters and related storage systems, configured in the same Site.

Note: If an associated provider is unregistered before, during, or after a Backup or Restore job executes, the job fails with a task framework error. If the unregistered providers are re-registered in IBM Spectrum Copy Data Management, new Backup or Restore jobs must be defined for the providers.

The following describes the procedures for registering a provider.

To register an Application Server - File System:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, select **Application Server**.
3. Right-click **Application Server**. Then click **Register** . The Register Application Server dialog opens.
4. Select **File System** as the Application Type.
5. Populate the fields in the dialog:

Site

A user-defined provider location, created in the **Sites & Providers**  view on the Configure tab.

Name

A user-defined name for the file system or volume. This can be the same as the host name or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique.

Host Address

A resolvable IP address or a resolvable path and machine name.

Port

The communications port of the provider you are adding. The default Windows port is 5985, and the default Linux and AIX port is 22,

OS Type

Select the file system or volume's operating system type. Available options include Windows, Linux, and AIX.

System Credential

Select or create your file system or volume's credentials. See [Identities Overview](#) on page 127.

6. Click **OK**. IBM Spectrum Copy Data Management first confirms a network connection and then adds the provider to the database.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

To register an Application Server - InterSystems Caché:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, select **Application Server**.
3. Right-click **Application Server**. Then click **Register** . The Register Application Server dialog opens.
4. Select **InterSystems Caché**  as the Application Type.
5. Populate the fields in the dialog:

Site

A user-defined provider location, created in the **Sites & Providers**  view on the Configure tab.

Name

A user-defined name for the InterSystems Caché server. This can be the same as the host name or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique.

Host Address

A resolvable IP address or a resolvable path and machine name.

Port

The communications port of the provider you are adding.

Run Inventory job after registration

If selected, IBM Spectrum Copy Data Management creates a high-level Inventory job and automatically catalogs the objects on the provider. Note that the Inventory job may take considerable time to complete.

Type

Select a **Virtual** or **Physical** InterSystems Caché server type. Select Virtual if the InterSystems Caché server is a VMware virtual machine.

If selecting Virtual, enter the vCenter location of the InterSystems Caché application server in the **vCenter** field.

OS Type

Select the InterSystems Caché server's operating system type. Available options include Windows, Linux, and AIX.

Authentication

IBM Spectrum Copy Data Management connects to the InterSystems Caché server as a local operating system user through an SSH key or password. See [Identities Overview](#) on page 127.

To use an SSH key, select **Key**, enter a username and select or create an SSH key.

To use a password, select **Password**, then select or create a Local credential.

System Credential:

Select or create your InterSystems Caché credentials. See [Identities Overview](#) on page 127.

6. Click **OK**. IBM Spectrum Copy Data Management first confirms a network connection and then adds the provider to the database.

To troubleshoot an application server after registration, use the Test & Configure option. This option verifies communication with the server, tests DNS settings between the IBM Spectrum Copy Data Management appliance and the server, and installs an IBM Spectrum Copy Data Management agent on the server. From the Provider Browser pane, right-click the application server, then click **Test & Configure** .

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

To register an Application Server - Oracle:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, select **Application Server**.
3. Right-click **Application Server**. Then click **Register** . The Register Application Server dialog opens.
4. Select **Oracle**  as the Application Type.
5. Populate the fields in the dialog:

Site

A user-defined provider location, created in the **Sites & Providers**  view on the Configure tab.

Name

A user-defined name for the Oracle server. This can be the same as the host name or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique.

Host Address

A resolvable IP address or a resolvable path and machine name. When registering an Oracle RAC cluster, register each node using its physical IP or name. Do not register a virtual name or SCAN (Single Client Access Name).

Run Inventory job after registration

If selected, IBM Spectrum Copy Data Management creates a high-level Inventory job and automatically catalogs the objects on the provider. Note that the Inventory job may take considerable time to complete.

Type

Select a **Virtual** or **Physical** Oracle server type. Select Virtual if the Oracle server is a VMware virtual machine. Note that AIX virtual servers should be registered as Physical.

If selecting Virtual, enter the vCenter location of the Oracle application server in the **vCenter** field.

Authentication

IBM Spectrum Copy Data Management connects to the Oracle server as a local operating system user through an SSH key or password. See [Identities Overview](#) on page 127.

To use an SSH key, select **Key**, enter a username and select or create an SSH key.

To use a password, select **Password**, then select or create a Local credential.

6. Click **OK**. IBM Spectrum Copy Data Management first confirms a network connection and then adds the provider to the database.

To troubleshoot an application server after registration, use the Test & Configure option. This option verifies communication with the server, tests DNS settings between the IBM Spectrum Copy Data Management appliance and the server, and installs an IBM Spectrum Copy Data Management agent on the server. From the Provider Browser pane, right-click the application server, then click **Test & Configure** .

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

To register an Application Server - SAP HANA:

Note: To successfully catalog a SAP HANA application provider, you must also register associated vCenters.

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, select **Application Server**.
3. Right-click **Application Server**. Then click **Register** . The Register Application Server dialog opens.
4. Select **SAP HANA**  as the Application Type.
5. Populate the fields in the dialog:

Site

A user-defined provider location, created in the **Sites & Providers**  view on the Configure tab.

Name

A user-defined name for the SAP HANA server. This can be the same as the host name or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique.

Host Address

A resolvable IP address or a resolvable path and machine name.

Port

The communications port of the provider you are adding. The format for the port number is 3<instance number>15. So, for example, if the instance number is 07, then enter the following port number: 30715.

Run Inventory job after registration

If selected, IBM Spectrum Copy Data Management creates a high-level Inventory job and automatically catalogs the objects on the provider. Note that the Inventory job may take considerable time to complete.

vCenter

The vCenter location of the SAP HANA application server.

System Credential / Database Credential(s):

Select or create your SAP HANA credentials. See [Identities Overview](#) on page 127.

6. Click **OK**. IBM Spectrum Copy Data Management first confirms a network connection and then adds the provider to the database.

To troubleshoot an application server after registration, use the Test & Configure option. This option verifies communication with the server, tests DNS settings between the IBM Spectrum Copy Data Management appliance and the server, and installs an IBM Spectrum Copy Data Management agent on the server. From the Provider Browser pane, right-click the application server, then click **Test & Configure** .

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

To register an Application Server - SQL:

When registering physical SQL servers it is recommended to register via the DNS server. The IBM Spectrum Copy Data Management appliance must be resolvable and route-able by the DNS server; the physical SQL server will communicate back to IBM Spectrum Copy Data Management through DNS.

Note that for physical SQL servers, you must allow outgoing connections to port 8443 on the IBM Spectrum Copy Data Management appliance from the SQL server.

Note: To successfully catalog a virtual SQL application provider, you must also register associated vCenters.

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, select **Application Server**.
3. Right-click **Application Server**. Then click **Register** . The Register Application Server dialog opens.
4. Select **SQL**  as the Application Type.
5. Populate the fields in the dialog:

Site

A user-defined provider location, created in the **Sites & Providers**  view on the Configure tab.

Name

A user-defined name for the SQL server. This can be the same as the host name or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique.

Host Address

A resolvable IP address or a resolvable path and machine name.

Port

The communications port of the provider you are adding. The default port is 5985.

In previous versions of IBM Spectrum Copy Data Management, the default SQL port was 1443.

SQL providers that were registered in previous versions using port 1443 will continue to function in newer versions of IBM Spectrum Copy Data Management.

Run Inventory job after registration

If selected, IBM Spectrum Copy Data Management creates a high-level Inventory job and automatically catalogs the objects on the provider. Note that the Inventory job may take considerable time to complete.

Type

Select a **Virtual** or **Physical** SQL server type. Select Virtual if the SQL server is a VMware virtual machine.

If selecting Virtual, enter the vCenter location of the SQL application server in the **vCenter** field.

vCenter

The vCenter location of the virtual SQL application server.

System Credential:

Select or create your SQL credentials. See [Identities Overview](#) on page 127.

Note: For Kerberos-based authentication only, the user identity must be specified in the username@FQDN format. The username must be able to authenticate using the registered password to obtain a ticket-granting ticket (TGT) from the key distribution center (KDC) on the domain specified by the fully qualified domain name.

6. Click **OK**. IBM Spectrum Copy Data Management first confirms a network connection and then adds the provider to the database.

To troubleshoot an application server after registration, use the Test & Configure option. This option verifies communication with the server, tests DNS settings between the IBM Spectrum Copy Data Management appliance and the server, and installs an IBM Spectrum Copy Data Management agent on the server. From the Provider Browser pane, right-click the application server, then click **Test & Configure** .

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

To register an Amazon Web Services (AWS) provider:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.

2. In the Provider Browser pane, select **AWS** .
3. Right-click **AWS** . Then click **Register** . The Register Amazon Web Services dialog opens.
4. Populate the fields in the dialog:

Site

A user-defined provider location, created in the **Sites & Providers**  view on the Configure tab.

Name

A user-defined name for the AWS provider. This can be the same as the host name or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique.

Region

Select an associated AWS region.

Comment

Optional provider description.

Access Key

Select or create your AWS access key. See [Identities Overview](#) on page 127.

5. Click **OK**. IBM Spectrum Copy Data Management first confirms a network connection and then adds the provider to the database.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

To register a DellEMC Unity provider:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, select **DellEMC Unity** .
3. Right-click **DellEMC Unity** . Then click **Register** . The Register dialog opens.
4. Populate the fields in the dialog:

Site

A user-defined provider location, created in the **Sites & Providers**  view on the Configure tab.

Name

A user-defined name for the DellEMC provider. This can be the same as the host name or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique.

Host Address

A resolvable IP address or a resolvable path and machine name.

Comment

Optional provider description.

Run Inventory job after registration

If selected, IBM Spectrum Copy Data Management creates a high-level Inventory job and automatically catalogs the objects on the provider. Note that the Inventory job may take considerable time to complete.

Credentials

Select or create your DellEMC Unity credentials. See [Identities Overview](#) on page 127.

Note: If upgrading from a previous version of IBM Spectrum Copy Data Management in which a username and password was entered during the provider registration process, an Identify will be automatically created for the provider.

5. Click **OK**. IBM Spectrum Copy Data Management first confirms a network connection and then adds the provider to the database.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

To register multiple DellEMC Unity providers through the Discover feature:

Use the Discover feature to find and register multiple DellEMC Unity providers by IP address or a range of IP addresses. For example: 172.27.*, 172.27.100.10-172.27.100.200, or 172.27.100.10.

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, select **DellEMC Unity** .
3. Right-click **DellEMC Unity** . Then click **Discover**. The Discover DellEMC Unity Providers dialog opens.
4. Enter an IP address or range of IP addresses associated with your **DellEMC Unity** storage systems in the **IP Address** field.
5. Click **Discover**. Discovered DellEMC Unity providers display.
6. Select providers to register along with universal custom parameters such as Site and credentials. To select individual parameters for each provider, click the parameters in the Custom Parameters field. Select specific Sites, credentials, ports and SSL parameters for each provider. If **Run Inventory job after registration** is selected, IBM Spectrum Copy Data Management creates a high-level Inventory job and automatically catalogs the objects on the provider.
7. Click **Register**. IBM Spectrum Copy Data Management adds the providers to the database.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

To register an IBM Spectrum Accelerate provider:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers**

- tab.
- In the Provider Browser pane, select **IBM Spectrum Accelerate** .
 - Right-click **IBM Spectrum Accelerate** . Then click **Register** . The Register dialog opens.
 - Populate the fields in the dialog:

Site

A user-defined provider location, created in the **Sites & Providers**  view on the Configure tab.

Name

A user-defined name for the IBM provider. This can be the same as the host name or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique.

Host Address

A resolvable IP address or a resolvable path and machine name.

Comment

Optional provider description.

Run Inventory job after registration

If selected, IBM Spectrum Copy Data Management creates a high-level Inventory job and automatically catalogs the objects on the provider. Note that the Inventory job may take considerable time to complete.

Credentials

Select or create your IBM Spectrum Accelerate credentials. See [Identities Overview](#) on page 127.

- Click **OK**. IBM Spectrum Copy Data Management first confirms a network connection and then adds the provider to the database.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

Note: IBM providers utilize port 22 for communication with IBM Spectrum Copy Data Management.

To register an IBM Spectrum Protect Snapshot provider:

Note: To successfully catalog an IBM Spectrum Protect Snapshot provider, you must also register associated vCenters and IBM storage providers.

- Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
- In the Provider Browser pane, select **IBM Spectrum Protect Snapshot** .
- Right-click **IBM Spectrum Protect Snapshot** . Then click **Register** . The Register dialog opens.
- Populate the fields in the dialog:

Site

A user-defined provider location, created in the **Sites & Providers**  view on the Configure tab.

Name

A user-defined name for the IBM provider. This can be the same as the host name or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique.

Host Address

A resolvable IP address or a resolvable path and machine name.

Comment

Optional provider description.

Run Inventory job after registration

If selected, IBM Spectrum Copy Data Management creates a high-level Inventory job and automatically catalogs the objects on the provider. Note that the Inventory job may take considerable time to complete.

Note: An automatic IBM Spectrum Protect Snapshot catalog job fails if associated IBM storage and vCenter providers are not registered. Ensure associated resources are registered before the automatic catalog job begins.

Credentials

Select or create your IBM Spectrum Protect credentials. See [Identities Overview](#) on page 127.

Note: When registering an IBM Spectrum Protect system, you must use credentials for the *tdpvmware* user.

Note: If upgrading from a previous version of IBM Spectrum Copy Data Management in which a username and password was entered during the provider registration process, an Identify will be automatically created for the provider.

5. Click **OK**. IBM Spectrum Copy Data Management first confirms a network connection and then adds the provider to the database.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

Note: IBM providers utilize port 22 for communication with IBM Spectrum Copy Data Management.

To register an IBM Spectrum Virtualize provider:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, select **IBM Spectrum Virtualize** .
3. Right-click **IBM Spectrum Virtualize** . Then click **Register** . The Register dialog opens.
4. Populate the fields in the dialog:

Site

A user-defined provider location, created in the **Sites & Providers**  view on the Configure tab.

Name

A user-defined name for the IBM provider. This can be the same as the host name or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique.

Host Address

A resolvable IP address or a resolvable path and machine name.

Comment

Optional provider description.

Run Inventory job after registration

If selected, IBM Spectrum Copy Data Management creates a high-level Inventory job and automatically catalogs the objects on the provider. Note that the Inventory job may take considerable time to complete.

Credentials

Select or create your IBM Spectrum Virtualize credentials. See [Identities Overview](#) on page 127.

Note: If upgrading from a previous version of IBM Spectrum Copy Data Management in which a username and password was entered during the provider registration process, an Identify will be automatically created for the provider.

5. Click **OK**. IBM Spectrum Copy Data Management first confirms a network connection and then adds the provider to the database.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

Note: IBM providers utilize port 22 for communication with IBM Spectrum Copy Data Management.

To register an LDAP provider:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, select **LDAP** .
3. Right-click **LDAP** . Then click **Register** . The Register LDAP Server dialog opens.
4. Populate the fields in the dialog:

Name

A user-defined name for the LDAP Server. Provider names must be unique.

Host Address

The IP address or resolvable logical node name of the LDAP server.

Port

The port on which the LDAP server is listening. The typical default port is 389 for non SSL connections or 636 for SSL connections.

Use SSL

Enable to establish a secure connection to the LDAP server.

Credentials

Select or create your LDAP credentials. See [Identities Overview](#) on page 127.

Note: If upgrading from a previous version of IBM Spectrum Copy Data Management in which a username and password was entered during the provider registration process, an Identify will be automatically created for the provider.

Base DN

The location where users and groups can be found.

User Filter

A filter to select only those users under the Base DN that match certain criteria. An example of a valid default user filter is **cn={0}**.

To enable authentication using the sAMAccountName Windows user naming attribute, set the User Filter to **samaccountname={0}**.

To enable authentication using an e-mail address associated with LDAP, set the User Filter to **mail={0}**.

Note that this entry also controls the type of user name that appears in IBM Spectrum Copy Data Management display of users.

User RDN

The relative distinguished path for the user. Specify the path where user records can be found. An example of a valid default RDN is:

cn=Users

Group RDN

The relative distinguished path for the group. Specify the path where group records can be found if the group is at a different level than the user path.

Comment

Optional description.

5. Click **OK**. IBM Spectrum Copy Data Management first confirms a network connection and then adds the provider to the database.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

To register a NetApp ONTAP provider:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers**

- tab.
- In the Provider Browser pane, select **NetApp ONTAP** .
 - Right-click **NetApp ONTAP** . Then click **Register** . The Register NetApp ONTAP dialog opens.
 - Populate the fields in the dialog:

Site

A user-defined provider location, created in the **Sites & Providers**  view on the Configure tab.

Name

A user-defined name for the NetApp storage system. This can be the same as the host name or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique.

Host Address

A resolvable IP address or a resolvable path and machine name.

Port

The communications port of the provider you are adding. Select the **Use SSL** check box to enable an encrypted Secure Socket Layer connection. The typical default port is 80 for non SSL connections or 443 for SSL connections.

Comment

Optional provider description.

Run Inventory job after registration

If selected, IBM Spectrum Copy Data Management creates a high-level Inventory job and automatically catalogs the objects on the provider. Note that the Inventory job may take considerable time to complete.

Credentials

Select or create your NetApp ONTAP credentials. See [Identities Overview](#) on page 127.

Note: If upgrading from a previous version of IBM Spectrum Copy Data Management in which a username and password was entered during the provider registration process, an Identify will be automatically created for the provider.

- Click **OK**. IBM Spectrum Copy Data Management first confirms a network connection and then adds the provider to the database.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

To register multiple NetApp ONTAP providers through the Discover feature:

Use the Discover feature to find and register multiple NetApp ONTAP providers by IP address or a range of IP addresses. For example: 172.27.*, 172.27.100.10-172.27.100.200, or 172.27.100.10.

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, select **NetApp ONTAP** .
3. Right-click **NetApp ONTAP** . Then click **Discover**. The Discover NetApp ONTAP Providers dialog opens.
4. Enter an IP address or range of IP addresses associated with your NetApp ONTAP storage systems in the **IP Address** field. Enter Simple Network Management Protocol settings such as the community name and protocol version in the **SNMP Options** fields. The default community name is "public."
5. Click **Discover**. Discovered NetApp ONTAP providers display.
6. Select providers to register along with universal custom parameters such as Site and credentials. To select individual parameters for each provider, click the parameters in the Custom Parameters field. Select specific Sites, credentials, ports and SSL parameters for each provider. If **Run Inventory job after registration** is selected, IBM Spectrum Copy Data Management creates a high-level Inventory job and automatically catalogs the objects on the provider.
7. Click **Register**. IBM Spectrum Copy Data Management adds the providers to the database.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

To register a HPE Nimble Storage provider:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, select **HPE Nimble Storage** .
3. Right-click **HPE Nimble Storage**. Then click **Register** . The Register dialog opens.
4. Populate the fields in the dialog:

Site

A user-defined provider location, created in the **Sites & Providers**  view on the Configure tab.

Name

A user-defined name for the HPE Nimble Storage provider. This can be the same as the host name or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique.

Host Address

A resolvable IP address or a resolvable path and machine name.

Port

The communications port of the provider you are adding. Select the **Use SSL** check box to enable an encrypted Secure Socket Layer connection. Enter the port number 5392. Select the **Use SSL** check box

to enable an encrypted Secure Socket Layer connection. The typical default port is 80 for non SSL connections or 443 for SSL connections.

Credentials

Select or create your HPE Nimble Storage credentials. See [Identities Overview](#) on page 127.

5. Click **OK**. IBM Spectrum Copy Data Management first confirms a network connection and then adds the provider to the database.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

To register a Pure Storage FlashArray provider:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, select **Pure Storage FlashArray** .
3. Right-click **Pure**. Then click **Register** . The Register dialog opens.
4. Populate the fields in the dialog:

Site

A user-defined provider location, created in the **Sites & Providers**  view on the Configure tab.

Name

A user-defined name for the Pure provider. This can be the same as the host name or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique.

Host Address

A resolvable IP address or a resolvable path and machine name.

Port

The communications port of the provider you are adding. Select the **Use SSL** check box to enable an encrypted Secure Socket Layer connection. The typical default port is 80 for non SSL connections or 443 for SSL connections.

Credentials

Select or create your Pure credentials. See [Identities Overview](#) on page 127.

5. Click **OK**. IBM Spectrum Copy Data Management first confirms a network connection and then adds the provider to the database.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

To register an SMTP provider:

Set up an SMTP server for email communications from IBM Spectrum Copy Data Management. Note that the SMTP provider is not cataloged.

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, select **SMTP** .
3. Right-click **SMTP** . Then click **Register** . The Register SMTP Server dialog opens.
4. Populate the fields in the dialog:

Name

A user-defined name for the SMTP server. This can be the same as the host name or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique.

Host Address

A resolvable IP address or a resolvable path and machine name.

Port

The communications port of the provider you are adding. Select the **Use SSL** check box to enable an encrypted Secure Socket Layer connection. The typical default port is 25 for non SSL connections or 443 for SSL connections.

Credentials

Select or create your SMTP credentials. See [Identities Overview](#) on page 127.

Note: If upgrading from a previous version of IBM Spectrum Copy Data Management in which a username and password was entered during the provider registration process, an Identify will be automatically created for the provider.

Comment

Optional provider description.

Click **Email Options** to set additional SMTP provider options:

From Address

Set the address to be associated with email communications from IBM Spectrum Copy Data Management.

Subject Prefix

Set a prefix to add to the email subject lines sent from IBM Spectrum Copy Data Management.

Timeout (msec)

Set the email timeout value in milliseconds.

5. Click **OK**. IBM Spectrum Copy Data Management first confirms a network connection and then adds the provider to the database.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

To register a VMware provider:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, select **VMware** .
3. Right-click **VMware** . Then click **Register** . The Register VMware Server dialog opens.
4. Populate the fields in the dialog:

Site

A user-defined provider location, created in the **Sites & Providers**  view on the Configure tab.

Name

A user-defined name for the VMware server. This can be the same as the host name or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique.

Host Address

A resolvable IP address or a resolvable path and machine name.

Port

The communications port of the provider you are adding. Select the **Use SSL** check box to enable an encrypted Secure Socket Layer connection. The typical default port is 80 for non SSL connections or 443 for SSL connections.

Credentials

Select or create your VMware credentials. See [Identities Overview](#) on page 127.

Note: If upgrading from a previous version of IBM Spectrum Copy Data Management in which a username and password was entered during the provider registration process, an Identify will be automatically created for the provider.

Comment

Optional provider description.

Run Inventory job after registration

If selected, IBM Spectrum Copy Data Management creates a high-level Inventory job and automatically catalogs the objects on the provider. Note that the Inventory job may take considerable time to complete.

5. Click **OK**. IBM Spectrum Copy Data Management first confirms a network connection and then adds the provider to the database.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

To add credentials to a registered virtual machine, see [Add Credentials to a Virtual Machine](#) on page **101**.

To register multiple VMware providers through the Discover feature:

Use the Discover feature to find and register multiple VMware providers by IP address or a range of IP addresses. For example: 172.27.*, 172.27.100.10-172.27.100.200, or 172.27.100.10.

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, select **VMware** .
3. Right-click **VMware** . Then click **Discover**. The Discover VMware Providers dialog opens.
4. Enter an IP address or range of IP addresses associated with your VMware providers in the **IP Address** field.
5. Click **Discover**. Discovered VMware providers display.
6. Select providers to register along with universal custom parameters such as Site and credentials. To select individual parameters for each provider, click the parameters in the Custom Parameters field. Select specific Sites, credentials, ports and SSL parameters for each provider. If **Run Inventory job after registration** is selected, IBM Spectrum Copy Data Management creates a high-level Inventory job and automatically catalogs the objects on the provider.
7. Click **Register**. IBM Spectrum Copy Data Management adds the providers to the database.

If a message appears indicating that the connection is unsuccessful, review your entries. If your entries are correct and the connection is unsuccessful, contact a system administrator to review the connections.

To add credentials to a registered virtual machine, see [Add Credentials to a Virtual Machine](#) on page **101**.

NEXT STEPS:

- Once providers are available in IBM Spectrum Copy Data Management and associated with a site, assign them to a resource pool. See [Configure Resource Pools](#) on page **106**.
- Add credentials to virtual machines in a VMware environment. See [Add Credentials to a Virtual Machine](#) on page **101**.

RELATED TOPICS:

- [Sites & Providers Overview](#) on page **70**
- [View a Provider](#) on page **95**
- [Edit a Provider](#) on page **98**
- [Unregister a Provider](#) on page **100**

- [Upload an SSL Certificate](#) on page **532**
- [LDAP User Name Syntax](#) on page **552**

View a Provider

Navigate through the Provider Browser to view a list of registered provider and resources that reside on those providers. The Provider Browser scans the actual provider and returns native properties.

BEFORE YOU BEGIN:

- Register providers before viewing them. See [Register a Provider](#) on page 75.

To view a list of providers:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the **Provider Browser** pane, expand the provider tree to view a list of all registered providers.
3. Drill down to view many of the objects, such as NetApp ONTAP volumes or VMware hosts, that reside on the providers.

To view job definitions associated with providers:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the **Provider Browser** pane, expand the provider tree to view a list of all registered providers.
3. Right-click a provider, then select **View Relationship**. All associated job definitions that interact with the provider display.

Use the Provider Details window to examine registered providers. The information presented depends on the type of provider selected.

LDAP Providers

Select an LDAP Server.

The LDAP Provider Details provides information about the LDAP server and the user accounts it manages.

The **General** tab provides information about the selected server including host address, port, and the use of SSL. The **Users** tab provides a list of all the user accounts configured on the server.

NetApp Providers

Select a volume within the NetApp ONTAP provider.

The NetApp ONTAP Provider Details provides information about the data storage in the selected volume. The **General** tab provides information about the selected volume including type, state, storage usage, reserve usage, and file usage. The gauges display the percentage of usage. Green represents data and blue represents

available space. The **Qtrees** tab lists qtrees stored in the selected volume and their status. The **Snapshots** tab lists the snapshots stored in the selected volume.

IBM Providers

Select a volume within the IBM provider.

The IBM Provider Details provides information about the data storage in the selected volume. The **General** tab provides information about the selected volume including the capacity, associated storage pool name, and mirrored copies synchronization rate.

DellEMC Providers

Select a volume within the DellEMC Unity provider.

The DellEMC Provider Details provides information about the data storage in the selected volume. The **General** tab provides information about the selected volume including the capacity, associated storage pool name, and mirrored copies synchronization rate.

SMTP Providers

Select an SMTP server.

The SMTP Provider Details provides information about the selected server including host address and port.

VMware Providers

Select a VMware vCenter within the VMware provider.

The **General** tab provides information about the selected vCenter including host address and software version. Use the **Hosts**, **VApps**, and **VMs** tabs to view a list of the virtual machine hosts, virtual appliances, and virtual machines that are configured on the selected vCenter.

The **Datacenters** tab provides a list of the datacenters configured on the vCenter. Use the **Datacenters** tab to view the list of **Datastores**, **Hosts**, and **VMs** filtered by the selected datacenter.

The **Datastores** tab lists the datastores configured on the vCenter. Select a datastore to view the list of **Hosts** and **VMs** filtered by the selected datastore.

Tip: Periodically closing tabs helps simplify navigation and browsing. To close multiple tabs, right-click a tab then select **Close Tab**, **Close Other Tabs**, or **Close All Tabs**.

Oracle Providers

Select an Oracle server.

The Oracle Provider Details provides information about the selected server including name and host address.

AWS Providers

Select an AWS provider.

The AWS Provider Details provides information about the selected server including region, name, type, and associated storage gateways.

NEXT STEPS:

- Run an Inventory job against storage providers that have not been recently cataloged. See [Jobs Overview](#) on page **161**.

RELATED TOPICS:

- [Sites & Providers Overview](#) on page **70**
- [Register a Provider](#) on page **75**
- [Edit a Provider](#) on page **98**
- [Unregister a Provider](#) on page **100**

Edit a Provider

Revise the properties of a provider as needed.

BEFORE YOU BEGIN:

- Review the properties of your current providers. See [Sites & Providers Overview](#) on page 70 for a list of supported providers.

To edit the properties of a provider:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, browse to the desired provider and select it.
3. Right-click the provider. Then click **Edit** . An update dialog opens.
4. Make revisions as needed. Fields to revise include the following:

Site

A user-defined provider location, created in the **Sites & Providers**  view on the Configure tab.

Name

A user-defined name for the provider. This can be the same as the host name or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique.

Host Address

A resolvable IP address or a resolvable path and machine name.

Port

The communications port of the provider you are adding. Select the **Use SSL** check box to enable an encrypted Secure Socket Layer connection. The typical default port is 80 for non SSL connections or 443 for SSL connections.

Credentials

Select or create your credentials. See [Identities Overview](#) on page 127.

Comment

Optional provider description.

5. Click **OK** when you are satisfied that the job-specific information is correct.

NEXT STEPS:

- If the storage provider you edited has not recently been cataloged, catalog it. See [Jobs Overview](#) on page **161**.

RELATED TOPICS:

- [Sites & Providers Overview](#) on page **70**
- [Register a Provider](#) on page **75**
- [View a Provider](#) on page **95**
- [Unregister a Provider](#) on page **100**

Unregister a Provider

Unregister a registered provider if you do not want to run reports against it, search for objects on it, or create a job.

Note: A provider cannot be deleted if it is assigned to a Resource Pool. Remove your providers from Resource Pools before deleting.

Best Practice: Ensure the provider you unregister is not associated with any defined job. To view job definitions associated with a provider, click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab. Right-click a provider from the Provider Browser view, then select **View Relationship**. All associated job definitions that interact with the provider display.

BEFORE YOU BEGIN:

- Review the properties of the provider to determine if you want to unregister it. See [Edit a Provider](#) on page 98.

Note: If an associated provider is unregistered before, during, or after a Backup or Restore job executes, the job fails with a task framework error. If the unregistered providers are re-registered in IBM Spectrum Copy Data Management, new Backup or Restore jobs must be defined for the providers.

To unregister a provider:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, browse to the desired provider and select it.
3. Right-click the provider. Then click **Unregister** . A confirmation dialog box opens.
4. Confirm unregistration. The provider is unregistered.

RELATED TOPICS:

- [Sites & Providers Overview](#) on page 70
- [Register a Provider](#) on page 75
- [View a Provider](#) on page 95
- [Edit a Provider](#) on page 98

Add Credentials to a Virtual Machine

Some features in IBM Spectrum Copy Data Management require credentials to access virtual machines in your VMware environment, such as truncating application logs when running a VMware Backup job. Credentials can be added to individual virtual machines or to multiple virtual machines if the credential information is universal.

BEFORE YOU BEGIN:

- At least one VMware provider must be registered and cataloged in IBM Spectrum Copy Data Management. See [Register a Provider](#) on page 75 and [Create an Inventory Job Definition - VMware](#) on page 194.

To add credentials for a single virtual machine:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, expand the **VMware** object and select a VMware provider.
3. Click the **VMs** tab to view associated virtual machines, then click the virtual machine name.
4. Click the **Credentials** tab.
5. Click **New** . The New Credential dialog opens.
6. Enter the username, password and an optional description in the **Comment** field.
7. Select the credential type. Options include System and SQL.
8. If entering SQL credentials, enter the name of the SQL instance in the **Instance Name** field.
9. To apply System credentials to application instances (for example, SQL instances), enable the **Use System Credentials for apps** option. Note that System credentials are always required. If your application instances use credentials that differ from your System credentials, you must repeat the above procedure for each application instance using different Instance Names.

To add credentials for multiple virtual machines:

1. Click the **Configure**  tab. On the Views pane, select **Sites & Providers** , then select the **Providers** tab.
2. In the Provider Browser pane, expand the **VMware** object and right-click a VMware provider. Then click **Manage VMs** . The Manage Virtual Machines dialog opens.
3. Enter a wildcard to search for virtual machines available on the VMware provider. For example, vm* or vm [1-50].
4. Select virtual machines with universal credentials.

5. Enter the universal credential information for the virtual machines, along with the credential type and instance name if applicable.
6. To apply System credentials to application instances (for example, SQL instances), enable the **Use System Credentials for apps** option. Note that System credentials are always required. If your application instances use credentials that differ from your System credentials, you must repeat the above procedure for each application instance.
7. Click **Close** to exit the Manage Virtual Machines dialog.

Existing credentials for multiple virtual machines can also be updated through the **Manage VMs**  feature. To update credentials for an individual virtual machine, click **Manage** in the row containing the virtual machine.

RELATED TOPICS:

- [Register a Provider](#) on page **75**
- [Create an Inventory Job Definition - VMware](#) on page **194**

Role-Based Access Control

The topics in the following section cover configuring Resource Pools, Roles, and Accounts.

Role-Based Access Control Overview

Role-based access control allows you to set the resources and permissions available to IBM Spectrum Copy Data Management accounts. Through role-based access control you can tailor IBM Spectrum Copy Data Management for individual users, giving them access to the features and providers they need. Once providers are associated with a site, they can be added to a resource pool along with high level IBM Spectrum Copy Data Management features such as Policies, Reports, and screens. Roles are then configured to define the actions that can be performed by the user of the account associated with the resource pool. These parameters are then associated with one or more user accounts, which can be native to IBM Spectrum Copy Data Management or imported as part of an LDAP group.

- [Configure Resource Pools](#) on page 106
- [Configure Roles](#) on page 109
- [Configure Accounts](#) on page 111
- [Configure Tenants](#) on page 122

Note: Users that register providers, such as storage devices, or add resources to IBM Spectrum Copy Data Management, such as jobs or customized reports, will have full access to interact with those providers or resources regardless of role-based access control restrictions. For example, if a user's permission allows them to register NetApp providers, they will also be able to view, edit, and unregister the NetApp providers that they registered, even if the necessary permissions are not assigned to them through role-based access control.

Configure role-based access control in the **Access Control**  view on the Configure tab.

Resource Pools - A resource pool defines the resources that will be made available to an account. Every provider added to IBM Spectrum Copy Data Management, such as storage devices and LDAP servers, can be included in a resource pool, along with individual IBM Spectrum Copy Data Management functions and screens. This gives you the ability to finely-tune the experience of a user. For example, a resource pool could include only storage devices associated with a single vendor, with access to only the IBM Spectrum Copy Data Management search and reporting functionality. When the resource pool is associated with a role and an account, the account user will only see the screens associated with search and reporting, and will only have access to the storage devices defined in the resource pool. See [Configure Resource Pools](#) on page 106.

Roles - Roles define the actions that can be performed on the resources defined in a resource pool. A resource pool defines the providers that will be made available to an account, such as storage devices, and resources, such as IBM Spectrum Copy Data Management functions and screens; a role sets the permissions to interact with the resources defined in the resource pool. For example, if a resource pool is created that includes IBM Spectrum Copy Data Management Backup and Restore jobs, the role will determine how a user can interact with the jobs. Permissions can be set to allow a user to create, view, and run the Backup and Restore jobs defined in a resource pool, but not delete them. Similarly, permissions can be set to create administrator accounts, allowing a user to create and edit other accounts, set up sites and resources, and interact with all of the available IBM Spectrum Copy Data Management features. See [Configure Roles](#) on page 109.

Accounts - An account associates a resource pool with a role. To enable a user to log on to IBM Spectrum Copy Data Management and use its functions, you must first add the user to IBM Spectrum Copy Data Management

as a native user or as part of an imported group of LDAP users, then assign resource pools and roles to the user account. The account will have access to the resources and features defined in the resource pool as well as the permissions to interact with the resources and features defined in the role. See [Configure Accounts](#) on page **111**.

RELATED TOPICS:

- [Configure Resource Pools](#) on page **106**
- [Configure Roles](#) on page **109**
- [Configure Accounts](#) on page **111**
- [VMware Admin Role-Based Access Control Configuration](#) on page **113**
- [NetApp ONTAP Admin Role-Based Access Control Configuration](#) on page **115**
- [IBM Admin Role-Based Access Control Configuration](#) on page **117**

Configure Resource Pools

A resource pool is a component of the role-based access system, and defines the resources that will be made available to an account. Every provider added to IBM Spectrum Copy Data Management, such as storage devices and LDAP servers, can be included in a resource pool, along with individual IBM Spectrum Copy Data Management functions and screens. This gives you the ability to finely-tune the experience of a user. For example, a resource pool could include only storage devices associated with a single vendor, with access to only the IBM Spectrum Copy Data Management search and reporting functionality. When the resource pool is associated with a role and an account, the account user will only see the screens associated with search and reporting, and will only have access to the storage devices defined in the resource pool.

Enhanced granularity is supported when configuring resource pools for VMware providers, allowing administrators to give permissions to users at the following levels: datastore, host, and virtual machine. Expand the datastore level to view folders containing all available datastores, hosts, and virtual machines on the datastore, then assign them to the resource pool as needed. Note that hosts are used as data destinations in Copy Data Management jobs, so you must ensure a user running these jobs has the correct permissions to interact with the hosts and jobs through role-based access control.

BEFORE YOU BEGIN:

- Create sites to assign to your providers. A site is a user-defined grouping of providers that is generally based on location. See [Add a Site](#) on page 72.
- Add providers to IBM Spectrum Copy Data Management and associate them with a site. See [Register a Provider](#) on page 75.

Add a Resource Pool

1. Click the **Configure**  tab. On the Views pane, select **Access Control** , then select the **Resource Pools** tab.
2. Click **New** . The New Resource Pool editor opens.
3. Click the **Step 1: Providers** tab. From the list of available sites and providers, select one or more providers to add to the resource pool. Note that your providers are grouped into sites, which allows you to add entire sites to the resource pool, or specific providers within the site. Individual storage virtual machines and VMware datacenters can also be selected for use with a resource pool.
4. Click the **Step 2: Jobs** tab. Select one or more job types, individual custom jobs, schedules, and scripts to include in the resource pool.
5. Click the **Step 3: Reports** tab. Select one or more report types or individual reports to include in the resource pool.
6. Click the **Step 4: Applications** tab. Select one or more application servers to include in the resource pool.
7. Click the **Step 5: Identities** tab. Select one or more keys and credentials to include in the resource pool.

8. Click the **Step 6: Access Control** tab. Select security options that will be configurable by accounts associated with this resource pool. Available options include All Roles, All Accounts, All Resource Pools, and All SLA Policies. For example, if **All Resource Pools** is selected in this step, users associated with this resource pool can create, view, edit, and delete Resource Pools, if paired with the necessary "resourcepool" permission, set on the **Roles** pane.
9. Click the **Step 7: Screens** tab. Select the IBM Spectrum Copy Data Management screens to include in the resource pool. For example, to create an account that does not have access to the Marketplace, exclude the Marketplace from the resource pool in this step.
10. Click the **Step 8: Finish** tab. Enter a name for your resource pool and a meaningful description. When you are satisfied that the entered information is correct, click **Finish**. The resource pool appears on the All Resources pane and can be applied to new and existing accounts.

Edit a Resource Pool

Revise a resource pool to change the selected resources and IBM Spectrum Copy Data Management features. Updated resource pool settings take affect once accounts associated with the resource pool log in to IBM Spectrum Copy Data Management.

Note: The *All Resources* resource pool cannot be edited.

1. Click the **Configure**  tab. On the Views pane, select **Access Control**  , then select the **Resource Pools** tab.
2. Select the resource pool to edit by clicking in the row containing the resource pool name.
3. Click **Edit**  . The Edit Resource Pool dialog opens.
4. Update the resources and IBM Spectrum Copy Data Management features to assign to the resource pool.
5. Click **Finish**. The revisions are applied to the resource pool.

Delete a Resource Pool

Delete a resource pool when it becomes obsolete.

A resource pool cannot be deleted if it is assigned to an account. On the All Resource Pools pane, click **View Relationship** to view the accounts that are associated with the resource pool. Re-assign your accounts to different resource pools before deleting.

Note: The *All Resources* resource pool cannot be deleted.

1. Click the **Configure**  tab. On the Views pane, select **Access Control**  , then select the **Resource Pools** tab.
2. Select the resource pool to delete by clicking in the row containing the resource pool name.
3. Click **Delete**  . A confirmation dialog box displays.
4. Confirm deletion. The resource pool is deleted.

NEXT STEPS:

- Create roles to define the actions that can be performed by the user of an account associated with a resource pool. Roles are used to define permissions to interact with the resources defined in the resource pool. See [Configure Roles](#) on page **109**.

RELATED TOPICS:

- [Configure Roles](#) on page **109**
- [Configure Accounts](#) on page **111**

Configure Roles

A role is a component of the role-based access system, and is used to define the actions that can be performed by the user of an account associated with a resource pool. A resource pool defines the resources that will be made available to an account, such as storage devices and IBM Spectrum Copy Data Management features; a role sets the permissions to interact with the resources defined in the resource pool. For example, if a resource pool is created that includes IBM Spectrum Copy Data Management Backup and Restore jobs, the role will determine how a user can interact with the jobs. Permissions can be set to allow a user to create, view, and run the Backup and Restore jobs defined in a resource pool, but not delete them. Similarly, permissions can be set to create administrator accounts, allowing a user to create and edit other accounts, set up sites and providers, and interact with all of the available IBM Spectrum Copy Data Management features.

BEFORE YOU BEGIN:

- Create sites to assign to your providers. A site is a user-defined grouping of providers that is generally based on location. See [Add a Site](#) on page 72.
- Add providers to IBM Spectrum Copy Data Management and associate them with a site. See [Register a Provider](#) on page 75.
- Once providers are available in IBM Spectrum Copy Data Management and associated with a site, assign them to a resource pool. See [Configure Resource Pools](#) on page 106.

Add a Role

1. Click the **Configure**  tab. On the Views pane, select **Access Control** , then select the **Roles** tab.
2. In the All Roles pane, click **New** . The New Role dialog opens.
3. Enter a role name and a meaningful description.
4. Select IBM Spectrum Copy Data Management features to add to the role, such as reports, jobs, and sites as well as provider types, such as VMware, LDAP, and SMTP.
5. When a feature is added to the role, it displays in the Permissions pane. Select permissions for the feature. For example, if the Site feature is added to the role, the following Site-based permissions are available: Create, View, Edit, Delete and All Permissions. If the Delete permission is excluded from the role, accounts associated with this role can create, view, and edit Sites, but cannot delete them.

Similarly, if the Report feature is added to the role, the Create permission allows accounts associated with the role to create custom reports. The View permission allows accounts associated with the role to view the list of reports in the Reports and Jobs tabs as well as run and view reports.

To set permissions for individual roles, click **Click to select permissions** next to the role name.

To set bulk permissions for multiple roles, select the check boxes next to the role names, then click the **Add Permissions**  drop-down menu. Select permissions to apply to the selected roles, then click **Apply**.

Permissions are then added to the selected roles.

6. When you are satisfied that the selected features and permissions are correct, click **Finish**. The role appears on the All Roles pane and can be applied to new and existing accounts.

Edit a Role

Revise a role to change the resources and permissions assigned to the role. Updated role settings take effect once accounts associated with the role log in to IBM Spectrum Copy Data Management.

Note: The SYSADMIN and USER roles cannot be edited.

1. Click the **Configure**  tab. On the Views pane, select **Access Control**  , then select the **Roles** tab.
2. In the All Roles pane, select the role to edit by clicking in the row containing the role name.
3. Click **Edit**  . The Edit Role dialog opens.
4. Select new resources and permissions to assign to the role.
5. Click **OK**. The revisions are applied to the role.

Delete a Role

Delete a role when it becomes obsolete.

A role cannot be deleted if it is assigned to an account. On the All Roles pane, click **View Relationship** to view the accounts that are associated with the role. Re-assign your accounts to different roles before deleting.

Note: The SYSADMIN, Read Only, and Create Only roles cannot be deleted.

1. Click the **Configure**  tab. On the Views pane, select **Access Control**  , then select the **Roles** tab.
2. In the All Roles pane, select the role to delete by clicking in the row containing the role name.
3. Click **Delete**  . A confirmation dialog box displays.
4. Confirm deletion. The role is deleted.

NEXT STEPS:

- Create an account. An account associates resource pools and roles with a user. Accounts can be native to IBM Spectrum Copy Data Management or can be imported as an LDAP group. See [Configure Accounts](#) on page 111.

RELATED TOPICS:

- [Configure Resource Pools](#) on page 106
- [Configure Accounts](#) on page 111

Configure Accounts

An account is a component of the role-based access system, and is used to associate resource pools and roles with a user. To enable a user to log on to IBM Spectrum Copy Data Management and use its functions, you must first add the user to IBM Spectrum Copy Data Management as a native user or as part of an imported group of LDAP users, then assign a resource pool and a role to the user account. The account will have access to the resources defined by the resource pool as well as the permissions to interact with the resources defined in the role.

Note that if multiple roles are assigned to a resource pool during account configuration, all permissions associated with the roles will be available to the account.

BEFORE YOU BEGIN:

- Create sites to assign to your providers. A site is a user-defined grouping of providers that is generally based on location. See [Add a Site](#) on page 72.
- Add providers to IBM Spectrum Copy Data Management and associate them with a site. See [Register a Provider](#) on page 75.
- Once providers are available in IBM Spectrum Copy Data Management and associated with a site, assign them to a resource pool. See [Configure Resource Pools](#) on page 106.
- Create roles to define the actions that can be performed by the user of an account associated with a resource pool. Roles are used to define permissions within a resource pool. See [Configure Roles](#) on page 109.

To add a native account to IBM Spectrum Copy Data Management:

1. Click the **Configure**  tab. On the Views pane, select **Access Control** , then select the **Accounts** tab.
2. Click **New** .
3. In the New Account pane, click **Create Native User**. The New Account dialog opens.
4. Enter a user name and password for the account.
5. Select one or more resource pools to add to the account.
6. Select roles to associate with each resource pool.
7. Click **Finish**. The account appears on the Accounts pane.

To import LDAP groups into IBM Spectrum Copy Data Management:

1. Click the **Configure**  tab. On the Views pane, select **Access Control** , then select the **Accounts** tab.
2. Click **New** .

3. In the New Account pane, click **Import LDAP Group**. The New Account dialog opens and a list of available LDAP groups displays.
4. Select one or more LDAP groups to assign to the selected account.
5. Select one or more resource pools to add to the account.
6. Select roles to associate with each resource pool.
7. Click **Finish**. The account appears on the Accounts pane.

Edit an Account

Revise an account to edit the username, password, associated resource pools and roles. Updated account settings take affect once the account logs in to IBM Spectrum Copy Data Management.

1. Click the **Configure**  tab. On the Views pane, select **Access Control**  , then select the **Accounts** tab.
2. In the Accounts pane, select the account to edit by clicking in the row containing the account name.
3. Click **Edit**. The Edit Role dialog opens.
4. Set a new username, password and select new resource pools and roles to assign to the account.
5. Click **OK**. The revisions are applied to the account.

Delete an Account

Delete an account to remove access to all IBM Spectrum Copy Data Management functions.

1. Click the **Configure**  tab. On the Views pane, select **Access Control**  , then select the **Accounts** tab.
2. In the Accounts pane, select the account to delete by clicking in the row containing the account name.
3. Click **Delete**. A confirmation dialog box displays.
4. Confirm deletion. The account is deleted.

NEXT STEPS:

- Ensure the user has access to the appropriate IBM Spectrum Copy Data Management resources as well as the necessary permissions to interact with the resources. See [Configure Resource Pools](#) on page 106 and [Configure Roles](#) on page 109.

RELATED TOPICS:

- [Configure Resource Pools](#) on page 106
- [Configure Roles](#) on page 109

VMware Admin Role-Based Access Control Configuration

The VMware Admin Role manages VM resources, runs Backup and Restore jobs on those resources, and generates VMware related reports. This user has full access to VMware resources, however does not have access to any storage resources. The VMware Admin will create the required Backup jobs using pre-defined storage workflows that have the necessary storage resources selected.

Additionally, the VMware administrator manages the IBM Spectrum Protect Snapshot resource and requires access to the IBM Spectrum Protect Snapshot resource, IBM storage, and the virtual machines that host the IBM Spectrum Protect Snapshot resource.

Resource Pool Configuration

Providers Tab

1. Set up the Providers screen to include the root level for all VMware resources, the IBM Spectrum Virtualize resources, and the IBM Spectrum Protect Snapshot resources.

Note: The VMware administrator must manage IBM Spectrum Protect Snapshot resources because access to VMware resources is required.

2. Select a specific SMTP server.

Jobs Tab

3. Select the root level of all VMware related jobs.
4. Select the root level of the IBM Spectrum Protect Snapshot Inventory job.
5. Select the root Report job.
6. Select the root All SLA Policies.
7. Select the root All Schedules.

Reports Tab

8. Under Protection Compliance, select the VMware related reports.
9. Under Storage Utilization, select the VMware related reports.

Access Control Tab

10. No Security Resources will be assigned to this role.

Screens Tab

11. Select all available Screens except Logs. The Logs function contains audit logs that this user should not have access to.

12. Name and submit the Resource Pool.

Role Configuration

13. Select the permissions listed above.

14. Name and submit the Role.

VMware Admin Account Configuration

15. Create a new account and link it to the newly created VMware Admin Resource Pool and Role.

RELATED TOPICS:

- [Configure Resource Pools](#) on page **106**
- [Configure Roles](#) on page **109**
- [Configure Accounts](#) on page **111**

NetApp ONTAP Admin Role-Based Access Control Configuration

The NetApp ONTAP Admin Role manages NetApp ONTAP resources, runs Backup and Restore jobs on those resources, and generates NetApp ONTAP related reports. The NetApp ONTAP Admin is also responsible for configuring work flow templates to allow other IBM Spectrum Copy Data Management users to run Backup and Restore jobs using NetApp ONTAP resources without having direct access to them.

Resource Pool Configuration

Providers Tab

1. Set up the Providers screen to include the root level for all NetApp resources.
2. Select a specific SMTP server.

Jobs Tab

3. Select the root level of all NetApp related jobs.
4. Select the root Report job.
5. Select the root All SLA Policies.
6. Select the root All Schedules.

Reports Tab

7. Select the root File Analytics tree.
8. Under Protection Compliance, select the NetApp RPO and NetApp Protection Usage.
9. Select the root Storage Protection tree.
10. Under Storage Utilization, select the NetApp related reports.

Access Control Tab

11. No Security Resources will be assigned to this role.

Screens Tab

12. Select all available Screens except Logs. The Logs function contains audit logs that this user should not have access to.
13. Name and submit the Resource Pool.

Role Configuration

14. Select the permissions listed above.
15. Name and submit the Role.

NetApp Admin Account Configuration

16. Create a new account and link it to the newly created NetApp Admin Resource Pool and Role.

RELATED TOPICS:

- [Configure Resource Pools](#) on page **106**
- [Configure Roles](#) on page **109**
- [Configure Accounts](#) on page **111**

IBM Admin Role-Based Access Control Configuration

The IBM Admin Role manages IBM resources, runs Backup and Restore jobs on those resources, and generates IBM related reports. The IBM Admin is also responsible for configuring work flow templates to allow other IBM Spectrum Copy Data Management users to run Backup and Restore jobs using IBM resources without having direct access to them.

Resource Pool Configuration

Providers Tab

1. Set up the Providers screen to include the root level for all IBM storage resources.
Note: IBM Spectrum Accelerate resources are managed by the VMware Admin Role.
2. Select a specific SMTP server.

Jobs Tab

3. Select the root level of IBM Inventory, IBM Backup, and IBM Restore jobs.
4. Select the root Report job.
5. Select the root All SLA Policies.
6. Select the root All Schedules.

Reports Tab

7. Under Protection Compliance, select the IBM RPO Compliance report.
8. Under Storage Utilization, select the IBM related reports.

Access Control Tab

9. No Security Resources will be assigned to this role.

Screens Tab

10. Select all available Screens except Logs. The Logs function contains audit logs that this user should not have access to.
11. Name and submit the Resource Pool.

Role Configuration

12. Select the permissions listed above.
13. Name and submit the Role.

IBM Admin Account Configuration

14. Create a new account and link it to the newly created IBM Admin Resource Pool and Role.

RELATED TOPICS:

- [Configure Resource Pools](#) on page 106
- [Configure Roles](#) on page 109
- [Configure Accounts](#) on page 111

HPE Nimble Storage Admin Role-Based Access Control Configuration

The HPE Nimble Storage Admin Role manages HPE Nimble Storage resources, runs Backup and Restore jobs on those resources, and generates HPE Nimble Storage related reports. The HPE Nimble Storage Admin is also responsible for configuring work flow templates to allow other IBM Spectrum Copy Data Management users to run Backup and Restore jobs using HPE Nimble Storage resources without having direct access to them.

Resource Pool Configuration

Providers Tab

1. Set up the Providers screen to include the root level for all HPE Nimble Storage resources.
2. Select a specific SMTP server.

Jobs Tab

3. Select the root level of all HPE Nimble Storage related jobs.
4. Select the root Report job.
5. Select the root All SLA Policies.
6. Select the root All Schedules.

Reports Tab

7. Under Protection Compliance, select the HPE Nimble Storage RPO Compliance.
8. Under Storage Utilization, select the HPE Nimble Storage related reports.

Access Control Tab

9. No Security Resources will be assigned to this role.

Screens Tab

10. Select all available Screens except Logs. The Logs function contains audit logs that this user should not have access to.
11. Name and submit the Resource Pool.

Role Configuration

12. Select the permissions listed above.
13. Name and submit the Role.

HPE Nimble Storage Admin Account Configuration

14. Create a new account and link it to the newly created HPE Nimble Storage Admin Resource Pool and Role.

RELATED TOPICS:

- [Configure Resource Pools](#) on page **106**
- [Configure Roles](#) on page **109**
- [Configure Accounts](#) on page **111**

Pure Storage FlashArray Admin Role-Based Access Control Configuration

The Pure Storage FlashArray Admin Role manages Pure Storage resources, runs Backup and Restore jobs on those resources, and generates Pure Storage related reports. The Pure Storage Admin is also responsible for configuring work flow templates to allow other IBM Spectrum Copy Data Management users to run Backup and Restore jobs using Pure Storage resources without having direct access to them.

Resource Pool Configuration

Providers Tab

1. Set up the Providers screen to include the root level for all Pure Storage resources.
2. Select a specific SMTP server.

Jobs Tab

3. Select the root level of all Pure Storage related jobs.
4. Select the root Report job.
5. Select the root All SLA Policies.
6. Select the root All Schedules.

Reports Tab

7. Under Protection Compliance, select the Pure Storage FlashArray RPO Compliance.
8. Under Storage Utilization, select the Pure Storage related reports.

Access Control Tab

9. No Security Resources will be assigned to this role.

Screens Tab

10. Select all available Screens except Logs. The Logs function contains audit logs that this user should not have access to.
11. Name and submit the Resource Pool.

Role Configuration

12. Select the permissions listed above.
13. Name and submit the Role.

Pure Storage Admin Account Configuration

14. Create a new account and link it to the newly created Pure Storage Admin Resource Pool and Role.

RELATED TOPICS:

- [Configure Resource Pools](#) on page **106**
- [Configure Roles](#) on page **109**
- [Configure Accounts](#) on page **111**

Configure Tenants

A tenant is a grouping of resources and users that are administered by a tenant administrator. An IBM Spectrum Copy Data Management administrator creates tenants, assigns resources to be made available to the tenants, and creates the tenant administrator. The tenant administrator can then further control and restrict resources for users in the tenant group, as well as add additional users to the tenant through LDAP. Tenants can be assigned shared resources, but in most cases would not have access to the resources or users of other tenants. Only IBM Spectrum Copy Data Management administrators and tenant administrators can configure a tenant; tenant users cannot configure a tenant.

A resource pool and a role determines the IBM Spectrum Copy Data Management resources and actions available within a tenant. A built-in Tenant role may be selected, which gives tenant users the ability to register resources, create job definitions, and other predefined IBM Spectrum Copy Data Management tasks.

To log in to the tenant, use the following format: tenant name/user name. For example, if the tenant is named "tenant1," a user with the username "tenant_user" would log in by entering the following in the IBM Spectrum Copy Data Management username field: tenant1/tenant_user.

To ensure tenant administrators and users can only view job definitions associated with their tenant, you must assign the Create permission, not the View permission, for jobs in the **Select the roles/permissions for the resource pool** step. Assigning the View permission gives tenant administrators and users full access to all jobs in the Resource Pool, including jobs that are not associated with the tenant. By granting only Create permissions for jobs, tenant administrators and users can create their own tenant-specific jobs. Tenant administrators can always view the jobs created by their tenant users, regardless of assigned permissions.

BEFORE YOU BEGIN:

- Review [Best Practices for Configuring Tenants](#) on page 124.
- Create a resource pool to associate with the tenant. A resource pool is a component of the role-based access system, and defines the resources that will be made available to the tenant. See [Configure Resource Pools](#) on page 106.
- Create a role to associate to the resource pool and the users of the tenant. A role defines the actions that can be performed on the resources defined in the tenant's resource pool. [Configure Roles](#) on page 109.

Add a Tenant

1. Click the **Configure**  tab. On the Views pane, select **Access Control** , then select the **Tenants** tab.
2. Click **New** . The New Tenant editor opens.
3. In the **Enter Tenant Info** section, enter a name for the tenant in the **Tenant Name** field as well as a tenant administrator name in the **Tenant Admin Name** field. Enter and confirm a password for the tenant administrator.
4. In the **Select resource pools** section, select one or more resource pools to add to the tenant.

5. In the **Select the roles/permissions for the resource pool** section, click **Click to select roles** to assign roles to the selected resource pools. Note that a built-in Tenant role may be selected, which gives tenant users the ability to register resources, create job definitions, and other predefined IBM Spectrum Copy Data Management tasks.
6. When you are satisfied that the entered information is correct, click **Finish**. The tenant appears on the All Tenants pane and the administrator account can log in to the newly created tenant using the following format: tenant name/tenant admin name.

Edit a Tenant

Revise a tenant to change the associated resource pools and permissions. Updated tenant settings take effect once accounts associated with the tenant log in.

1. Click the **Configure**  tab. On the Views pane, select **Access Control** , then select the **Tenants** tab.
2. Select the tenant to edit by clicking in the row containing the tenant name.
3. Click **Edit** . The Update Tenant Info editor displays.
4. Update the Tenant Name, Tenant Admin Name, and resource pools associated with the tenant.
5. Click **Finish**. The revisions are applied to the tenant.

Delete a Tenant

Delete a tenant when it becomes obsolete. Note that before deletion, associated jobs and resources must be cleaned up through the Maintenance job. The Maintenance job removes resources and associated objects created by IBM Spectrum Copy Data Management when a job in a pending state is deleted. The cleanup procedure reclaims space on your storage devices, cleans up your IBM Spectrum Copy Data Management catalog, and removes related snapshots.

1. Click the **Configure**  tab. On the Views pane, select **Access Control** , then select the **Tenants** tab.
2. Select the tenant to delete by clicking in the row containing the tenant name.
3. Click **Delete** . A confirmation dialog box displays.
4. Confirm deletion. The tenant is deleted.

RELATED TOPICS:

- [Best Practices for Configuring Tenants](#) on page 124
- [Configure Resource Pools](#) on page 106

Best Practices for Configuring Tenants

Review the following best practices when creating new tenants and assigning roles for specific use cases.

To assign resources to a tenant without granting the tenant users the ability to modify or delete the resources:

Create a new resource pool, and assign resources to be made available to the tenant in a resource pool on the **Resource Pools** tab. On the **Tenants** tab, select the newly created resource pool and assign the Read Only permission. This allows a tenant user to view the resources defined in the resource pool, but not modify or delete them.

To assign permissions to a tenant that allows tenant users to create new job definitions and reports, but prevents them from viewing existing IBM Spectrum Copy Data Management job definitions:

Create a new resource pool, and assign resources to be made available to the tenant in a resource pool on the **Resource Pools** tab. On the **Tenants** tab, select the newly created resource pool and assign the Create Only permission. This allows a tenant user to create new job definitions and reports, but prevents them from viewing existing IBM Spectrum Copy Data Management job definitions.

To assign resources to a tenant and allow the tenant users to create and run job sessions, reports and perform searches:

Create a new resource pool, and assign resources to be made available to the tenant in a resource pool on the **Resource Pools** tab. On the **Tenants** tab, select the newly created resource pool and assign the Read Only and Create Only permissions. This allows a tenant user to create as well as run new job sessions and reports.

General Recommendations

- For a tenant admin, it is recommended to create two resource pools. In the first resource pool, add the providers to be made available to the tenant, and assign a Read Only permission to the resource pool. In the second resource pool, assign jobs, security, and screens, and assign the Create Only permission. Once complete, assign both resource pools to the tenant.
- When configuring a resource pool for a tenant user, it is recommended to exclude the security resources found in the **Step 5: Security** step or the **Home** and **Logs** resources found in the **6. Screens** step. These resources contain general IBM Spectrum Copy Data Management information that may not apply to the tenant user.
- The built-in *All Resources* resource pool should not be assigned to a tenant as it includes all of the resources in the IBM Spectrum Copy Data Management system.
- Selecting higher level objects instead of specific resources and assigning the View, Edit, and Delete permissions may cause tenants to see resources from other tenants. Add lower-level resources to ensure the tenants can only see objects assigned to the tenant.

RELATED TOPICS:

- [Configure Tenants](#) on page **122**
- [Configure Resource Pools](#) on page **106**

Identities

The topics in the following section cover adding SSH keys and adding, editing, and deleting credentials.

Identities Overview

Some features in IBM Spectrum Copy Data Management require credentials and keys to access your providers. For example, IBM Spectrum Copy Data Management connects to the Oracle servers as the local operating system user specified during registration in order to perform tasks like cataloging, data protection, and data restores. IBM Spectrum Copy Data Management also logs into local database and ASM instances as this user through password-less OS authentication. Therefore, the user must have all the privileges IBM Spectrum Copy Data Management needs to perform its tasks.

Credentials and keys are configured through the **Identities**  view.

RELATED TOPICS:

- [Add a Key](#) on page **128**
- [Add a Credential](#) on page **132**

Add a Key

Some features in IBM Spectrum Copy Data Management require credentials and keys to access your providers. For example, IBM Spectrum Copy Data Management connects to the Oracle servers as the local operating system user specified during registration in order to perform tasks like cataloging, data protection, and data restores. IBM Spectrum Copy Data Management also logs into local database and ASM instances as this user through password-less OS authentication. Therefore, the user must have all the privileges IBM Spectrum Copy Data Management needs to perform its tasks.

IBM Spectrum Copy Data Management connects to Oracle servers as a local operating system user through a password or an SSH key. To use a key, enter a username and select or create an SSH key. When using a key, the username must exist as a local user on the Oracle server. For password-based authentication, the password must be correctly configured for the appropriate user on the Oracle server. For key-based authentication, the public key must be placed in the `authorized_keys` file for the appropriate user on the Oracle server. See [Oracle Requirements](#) on page 41.

Amazon Web Services (AWS) access keys and secret keys are configured through the AWS Management Console and then added to IBM Spectrum Copy Data Management.

The procedures below describe how to add keys and register associated Oracle or AWS providers.

Add an SSH key through the *Generate a keypair for me* method and register an associated provider

1. In IBM Spectrum Copy Data Management, click the **Configure**  tab. On the Views pane, select **Identities** , then the **Keys** tab.
2. Click **New** . The Create Key dialog displays.
3. Select **SSH** as the key type and enter a key name in the Name field.
4. Select **Generate a keypair for me** as the creation type and enter an optional comment. Click **OK**. A public key is generated and displays in the Create Key dialog. Copy the key. See the following steps to use this key to register an Oracle provider.
5. On the Oracle server, execute `cd ~/ . ssh` while logged in as Oracle user assigned to IBM Spectrum Copy Data Management. Paste and save the generated public key to the `authorized_keys` file.
6. In IBM Spectrum Copy Data Management, click the **Configure**  tab. On the Views pane, select **Sites & Providers** . The Provider Browser opens.
7. Right-click **Oracle** in the Provider Browser, then click **Register** . The Register Oracle Server dialog opens.
8. Select a Site, enter a Name and Host Address.
9. Select **Key** as the Authentication type. Enter the Oracle username, then select the key created in Step 2 in the Key field. Click OK.

Add an SSH key through the *I will provide a keypair* method and register an associated provider

Generation of keys can occur on the IBM Spectrum Copy Data Management appliance using the command-line interface (CLI) or any other compatible server. In some circumstances, creating and adding a private/public SSH keypair generated on another host may be desirable. It is possible to generate SSH keypairs on another computer and then import them onto the IBM Spectrum Copy Data Management appliance as needed.

Note: Generally, private keys should not be generated on a client server and then transferred to the IBM Spectrum Copy Data Management appliance. It is strongly suggested that appropriate security measures be taken to protect the secrecy of the private keys. Loss or exposure of SSH private keys outside of the SSH host can severely compromise the security of communications using the SSH protocol. It is not recommended to copy private keys between different systems. If a new SSH keypair is needed, it is strongly advised that the procedure in the **Add an SSH key through the Generate a keypair for me method and register an associated provider** topic be followed to have the IBM Spectrum Copy Data Management appliance generate the keypair and then copy the public key to the intended host. **If there is a special need to generate a keypair on another host, use the procedure outlined below and ensure that appropriate security measures are taken to create, secure, and enter the private key.**

1. Identify a machine that has SSH installed. This machine will be used to generate the new SSH keypair. Log in to the identified machine and launch the terminal.
2. In the terminal, generate an SSH keypair by using the `ssh-keygen` command. Execute the `ssh-keygen` command:

```
$ ssh-keygen
```

3. When prompted, enter the full path name where the key pair will be output. A default file will be suggested by the `ssh-keygen` command. The default should only be used if a key has not yet been generated, otherwise, using the default may overwrite an existing SSH key pair. The default will typically appear as `/home/<user_account>/.ssh/id_rsa.pub` where `<user_account>` is the account used to log in to this system. Any valid path name could be used for the new SSH key, for example `/home/<user_account>/newkey`. If a key with the default name already exists, this will be indicated with the message displayed below. Be careful not to overwrite preexisting keys if they are in use and only overwrite these files if you intend to do so. Press **N** to enter a different file in which to save the key to avoid overwriting an existing keypair.

```
/home/<user_account>/.ssh/id_rsa already exists. Overwrite (y/n)?
```

4. Supply a passphrase and press **Enter**. Otherwise, press **Enter** for no passphrase.
5. If a passphrase was supplied, enter it again. Press **Enter**.
6. The key generation will produce two files, one with the path name supplied in the previous steps for the private key, and another ending in `.pub` is the public key. Using the default naming, this will be `id_rsa` and `id_rsa.pub`. The generated public key (`id_rsa.pub`) will need to be transferred to the server to which the IBM Spectrum Copy Data Management appliance will connect. In this example, it will be an Oracle server. Transfer the public key to the Oracle server. For the remainder of this procedure, it is assumed that the keypair is saved in the default location using the default file names for the keypair: `/home/<user_account>/.ssh/`. If the keypair is created using a different file name, use that file name in the steps that follow.
7. On the server to which the IBM Spectrum Copy Data Management appliance will connect and to which the

public key has been copied, the key (`ida_rsa.pub`) will need to be appended to the user's `authorized_keys` file. The `authorized_keys` file is generally found in the user's SSH directory. For example, it may be found at the following location: `/home/<user_account>/.ssh/authorized_keys`. If the `authorized_keys` file does not exist, consult the operating system's documentation for the procedure to properly creating this file. If the file exists, append the contents of the public key to the `authorized_keys` file. If this is not being done from the account that contains the `authorized_keys` file, it may be necessary to enter the `su` command to switch to that user. The step below assumes that you are logged into the server with the account that contains the `authorized_keys` file:

```
$ cat ida_rsa.pub >> authorized_keys
```

Note: This process can be automated using the `ssh-copykey` program from the computer used to generate the key. Consult the vendor's documentation for details on usage of this program.

8. Log in to the IBM Spectrum Copy Data Management appliance.
9. Click on the **Configure**  tab. On the **Views** pane, select **Identities**  , and then click on the **Keys** tab.
10. Click **New**  . The Create Key dialog displays.
11. Select **SSH** as the key type and enter a name for the key in the **Name** field.
12. Select **I will provide a key pair** as the creation method.
13. On the server where the SSH keypair was generated, locate the private key (`ida_rsa`). For example, the key generated by this process is in the following directory: `/home/<user_account>/.ssh/`. Copy the contents of the private key (`ida_rsa`) to the IBM Spectrum Copy Data Management appliance into the **Private Key** field in the **Create Key** dialog.
14. **(Optional)** It is highly recommended to copy the public key (`ida_rsa.pub`) into the **Public Key** field.
15. **(Optional)** Enter a helpful comment so that the usage of the key can be easily recalled.
16. Click **OK** to create the key.
17. Once the key has been added to the IBM Spectrum Copy Data Management appliance, the server to which the IBM Spectrum Copy Data Management appliance will connect needs to be registered. In this example, an Oracle server is used. In the IBM Spectrum Copy Data Management appliance, click on the **Configure** tab.
18. On the **Views** pane, select **Sites & Providers**  . The Provider Browser opens.
19. Right-click **Oracle** in the **Provider Browser** dialog and then click **Register**  . The Register Oracle Server dialog opens.
20. Select a **Site**, enter a name in the **Name** field and a host address in the **Host Address** field.
21. Select **Key** as the Authentication type. Enter the username of the user account to which the public key was appended to the `authorized_keys` file on the host to which the IBM Spectrum Copy Data Management appliance will connect in Step 7. In this example, it is the Oracle server. Enter the **Oracle username**.
22. Select the key created in Step 10 in the **Key** field. Click **OK**.

Add an Amazon Web Services (AWS) key and register an associated provider

1. Create your AWS access key and secret key through the AWS Management Console. Make note of the access and secret keys, which will be used later in this procedure. See [Managing Access Keys for IAM Users](#).
2. In IBM Spectrum Copy Data Management, click the **Configure**  tab. On the Views pane, select **Identities** , then the **Keys** tab.
3. Click **New** . The Create Key dialog displays.
4. Select **AWS** as the key type and enter a key name in the Name field.
5. Enter the access key and secret key created in Step 1 in the Access and Secret fields. Enter an optional comment. See the following steps to use this key to register an AWS provider.
6. In IBM Spectrum Copy Data Management, click the **Configure**  tab. On the Views pane, select **Sites & Providers** . The Provider Browser opens.
7. Right-click **AWS**  in the Provider Browser, then click **Register** . The Register Amazon Web Services dialog opens.
8. Select a Site, enter a Name and select a Region.
9. Click **Select** in the Access Key section, then select the key created above. Click **OK**.

RELATED TOPICS:

- [Add a Credential](#) on page **132**
- [Oracle Requirements](#) on page **41**
- [Register a Provider](#) on page **75**

Add a Credential

Add a credential

1. Click the **Configure**  tab. On the Views pane, select **Identities** , then the **Credentials** tab.
2. Click **New** . The Create Credential dialog opens.
3. Select a credential type in the Type field. Available options include **System** and **Oracle**.
4. Enter a name for the credential in the Name field.
5. Enter your login information for the associated provider in the Username and Password fields. For example, if creating a credential for an Oracle database, enter your login information associated with the Oracle database.
6. Enter an optional comment, then click **OK**. The credential appears on the **Credentials** pane and can be applied to new and existing storage providers.

Edit a Credential

Revise a credential to change the associated username and password.

1. Click the **Configure**  tab. On the Views pane, select **Identities** , then the **Credentials** tab.
2. In the Credentials pane, select the credential to edit by clicking in the row containing the credential name.
3. Click **Edit** . The Edit Credential dialog opens.
4. Update the name, username, and password assigned to the credential.
5. Click **OK**. The revisions are applied to the credential.

Delete a Credential

1. Click the **Configure**  tab. On the Views pane, select **Identities** , then the **Credentials** tab.
2. In the Credentials pane, select the credential to delete by clicking in the row containing the credential name.
3. Click **Delete** . A confirmation dialog box displays.
4. Confirm deletion. The credential is deleted.

RELATED TOPICS:

- [Identities Overview](#) on page 127
- [Register a Provider](#) on page 75

Configure SLA Policies

SLA Policies allow storage and virtualization administrators to create customized templates for the key processes involved in the creation and use of Backup jobs. Copy types, destinations, and parameters are configured in SLA Policies, which can be used and re-used in Backup jobs.

Generally, a storage administrator creates SLA Policies after registering storage providers in IBM Spectrum Copy Data Management and creating accounts that will create, edit, and run Backup and Restore jobs through role-based access control. When configuring a Backup job definition, available SLA Policies display in the job creation wizard, tailored to the type of Backup job being created. VMware Backup jobs support DellEMC Unity, IBM, and NetApp SLA Policies.

IBM Spectrum Copy Data Management Backup jobs utilizing a VM Replication-based SLA Policy can also be directed to Amazon Web Services (AWS) S3 destinations through an AWS Storage Gateway.

In a NetApp ONTAP-based SLA Policy, after an initial primary snapshot is added to the job, additional vaults and mirrors ensure your data is replicated to multiple locations.

BEFORE YOU BEGIN:

- Register storage providers to be used in the SLA Policies. Assign the storage providers to Sites. See [Register a Provider](#) on page 75 and [Add a Site](#) on page 72.
- If configuring an AWS-based SLA Policy, ensure the AWS provider has been correctly added to IBM Spectrum Copy Data Management. AWS access keys and secret keys must first be configured through the AWS Management Console and then added to IBM Spectrum Copy Data Management. See the "Add an Amazon Web Services (AWS) key and register an associated provider" section in [Add a Key](#) on page 128.
- Create an account with the necessary permissions to create and run Backup jobs. See [Role-Based Access Control Overview](#) on page 104.
- All related NetApp ONTAP storage resources associated with a VMware provider must be added to IBM Spectrum Copy Data Management, which include NetApp ONTAP storage controllers and clusters. See [Register a Provider](#) on page 75.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- Note that a VADP-based VM Replication of a virtual machine with vRDM will convert the vRDM to a VMDK. When restoring the virtual machine, consider the size requirements of the virtual machine in addition to the vRDM when selecting a destination datastore. The original datastore may not have enough free space to store the converted vRDM file.

Physical RDMs are not supported.

- Note that VMware Backup and Restore jobs only support vCenters or ESX hosts running vSphere 5.1 through 6.5.
- Restoring to an older snapshot is not supported if a new snapshot is being used for mirror replication.

CONSIDERATIONS FOR VMWARE VIRTUAL VOLUMES:

- SLA Policies that include virtual machines stored on virtual volume (VVOL) datastores through VM Replication sub-policies are supported. Replication is supported on the VM Replication target.
- Storage snapshots of virtual machines that reside on a VVOL are currently not supported. If a storage snapshot operation is selected for a virtual machine that resides on a VVOL, the virtual machine is skipped.

To create an SLA Policy:

1. Click the **Configure**  tab. On the Views pane, select **SLA Policies** . The All SLA Policies pane opens.
2. In the All SLA Policies pane, click **New** . The New SLA Policies pane opens.
3. Select a type of policy to create based on your storage provider. Select **AWS** to create an Amazon Web Services Backup job containing VM copies, **NetApp ONTAP** to create a NetApp ONTAP Backup job containing snapshots, VM copies, mirrors and vaults, **IBM** to create an IBM Backup job containing FlashCopies, Global Mirrors with Change Volumes, and VM Copies, or **DellEMC Unity** to create a DellEMC Unity Backup job containing snapshots and replication copies. VMware Backup jobs support IBM, DellEMC Unity, and NetApp ONTAP SLA Policies depending on your storage provider.
4. Enter a name and a meaningful description of the SLA Policy.

Configure AWS SLA Policies

To add a VM Replication sub-policy to an AWS SLA Policy:

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add VM Copy** .
3. In the VM Replication Destination pane select an AWS cloud destination from the list of available resources as the VM Replication destination, along with an associated storage gateway. Note that storage gateways are discovered based on the region of the AWS provider.

4. In the Options pane set the VM Replication sub-policy options.

Keep Copies

After a certain number of copies are created for a resource, older copies are purged from the gateway. Enter the age of the copies to purge in the **Days** field, or the number of copies to keep in the **Copies (maximum)** field.

Target Volume Prefix Label

Enter an optional label to identify the target volume. This label is added as a prefix to the volume name created by the job.

Note: Volume prefix labels must contain only alphanumeric characters and underscores. Labels cannot begin with numeric characters.

Configure DellEMC Unity SLA Policies

To add a VM Replication sub-policy to a DellEMC Unity SLA Policy:

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add VM Copy** .
3. In the VM Replication Destination pane select a DellEMC host destination from the list of available resources as the VM Replication destination, along with an associated storage pool. If no storage pool is selected, the storage pool with the largest amount of space available is chosen by default.
4. In the Options pane set the VM Replication sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Note: Snapshot labels must contain only alphanumeric characters and underscores.

Name

Enter an optional label to replace the default snapshot sub-policy label displayed in IBM Spectrum Copy Data Management. The default initial label is VM Copy0

Access Type

Select the access type for file-based storage. Available access types include **Hidden .ckpt folder (read-only)** and **Shares**, based on CIFS or NFS mounting.

Destination storage limit in GB / Destination volumes limit

Specify quotas for storage usage and the number of volume created on the destination for all jobs utilizing the SLA Policy.

To add a snapshot sub-policy to a DellEMC Unity SLA Policy:

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add Snapshot** . In the Options pane, set the snapshot sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Name

Enter an optional name to replace the default snapshot sub-policy name displayed in IBM Spectrum Copy Data Management. The default initial name is Snapshot0.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Note: Snapshot labels must contain only alphanumeric characters and underscores.

Access Type

Select the access type for file-based storage. Available access types include **Hidden .ckpt folder (read-only)** and **Shares**, based on CIFS or NFS mounting.

To add a Replication sub-policy to a DellEMC Unity SLA Policy:

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add Replication** .
3. In the Replication Destination pane select a DellEMC host destination from the list of available resources as the Replication destination, along with an associated storage pool. If no storage pool is selected, the storage pool with the largest amount of space available is chosen by default.
4. In the Options pane set the Replication sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Name

Enter an optional name to replace the default Replication sub-policy name displayed in IBM Spectrum Copy Data Management. The default initial name is Replication0.

Keep Source Volume name for target volume

Enable to retain the source volume name for copy data generated by IBM Spectrum Copy Data Management.

Target Volume Prefix Label

Enter an optional label to identify the target volume. This label is added as a prefix to the volume name created by the job.

Note: Volume prefix labels must contain only alphanumeric characters and underscores. Labels cannot begin with numeric characters.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Note: Snapshot labels must contain only alphanumeric characters and underscores.

Synchronization

Specify the time in minutes in which the change volumes will be synchronized with a consistent copy of the data. If a copy does not complete in the cycle period, the next cycle period will not start until the copy is complete. Synchronization can also be initiated manually.

Destination storage limit in GB / Destination volumes limit

Specify quotas for storage usage and the number of volume created on the destination for all jobs utilizing the SLA Policy.

Configure IBM Spectrum Virtualize SLA Policies

To add a VM Replication sub-policy to an IBM Spectrum Virtualize SLA Policy:

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add VM Copy** .
3. In the Destination pane select an IBM host destination from the list of available resources as the VM Replication destination, along with an associated storage pool. If no storage pool is selected, the storage pool with the largest amount of space available is chosen by default. To select the original target destination, select **Use Original**.
4. In the Options pane set the VM Replication sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Target Volume Prefix Label

Enter an optional label to identify the target volume. This label is added as a prefix to the volume name created by the job.

Note: Volume prefix labels must contain only alphanumeric characters and underscores. Labels cannot begin with numeric characters.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Note: Snapshot labels must contain only alphanumeric characters and underscores.

Name

Enter an optional label to replace the default snapshot sub-policy label displayed in IBM Spectrum Copy Data Management. The default initial label is VM Copy0.

Protocol

If more than one storage protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

Full Copy Method

Select the full copy method. Available full copy methods include Clone or VADP-based VM Replication.

Destination storage limit in GB / Destination volumes limit

Specify quotas for storage usage and the number of volume created on the destination for all jobs utilizing the SLA Policy.

To add a FlashCopy sub-policy to an IBM Spectrum Virtualize SLA Policy:

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. In the Target FlashCopy Storage Pool pane select an IBM host destination from the list of available resources as the FlashCopy destination, along with an associated storage pool. If no storage pool is selected, the storage pool with the largest amount of space available is chosen by default. To select the original target destination, select **Use Original**. Note that you can select multiple storage pools from multiple IBM resources, but only one storage pool is allowed for each node.
3. To enable an Incremental FlashCopy, select **Enable Incremental FlashCopy** in the Incremental FlashCopy Storage Pool pane. Select an IBM host destination from the list of available resources as the FlashCopy destination, along with an associated storage pool. If no storage pool is selected, the storage pool with the largest amount of space available is chosen by default. To select the original target destination, select **Use Original**.

If the Enable Incremental FlashCopy option is selected, note that the base FlashCopy will be sent to the destination selected in the Incremental FlashCopy Storage Pool pane. Subsequent incremental FlashCopies will be sent to the destination selected in the Target FlashCopy Storage Pool pane.

Note that the Target FlashCopy Storage Pool must reside on the same storage system as the Incremental FlashCopy Storage Pool.

Note: When editing an existing SLA Policy, Incremental FlashCopy options cannot be altered.

4. Click **Add FlashCopy** . In the Options pane, set the FlashCopy sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Name

Enter an optional name to replace the default FlashCopy sub-policy name displayed in IBM Spectrum Copy Data Management. The default initial name is FlashCopy0.

FlashCopy Volume Prefix

Enter an optional label to identify the FlashCopy. This label is added as a prefix to the FlashCopy name created by the job.

Note: FlashCopy labels must contain only alphanumeric characters and underscores.

To add a Global Mirror with Change Volumes sub-policy to an IBM Spectrum Virtualize SLA Policy:

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add Global Mirror with Change Volumes**  .
3. In the Global Mirror with Change Volumes Destination pane select an IBM host destination from the list of available resources as the Global Mirror destination, along with an associated storage pool. If no storage pool is selected, the storage pool with the largest amount of space available is chosen by default. To select the original target destination, select **Use Original**.
4. In the Options pane set the Global Mirror with Change Volumes sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Name

Enter an optional name to replace the default Global Mirror sub-policy name displayed in IBM Spectrum Copy Data Management. The default initial name is Global Mirror0.

Keep Source Volume name for target volume

Enable to retain the source volume name for copy data generated by IBM Spectrum Copy Data Management.

Volume Prefix Label

Enter an optional label to identify the volume. This label is added as a prefix to the volume name created by the job and cannot be edited after the job is submitted.

Note: Volume prefix labels must contain only alphanumeric characters and underscores. Labels cannot begin with numeric characters.

Cycle Period (seconds)

Specify the time in which the change volumes will be refreshed with a consistent copy of the data. If a copy does not complete in the cycle period, the next cycle period will not start until the copy is complete. The range of possible values is 60 through 86400. The default is 300.

Global Mirror Volume Prefix

Enter an optional label to identify the Global Mirror. This label is added as a prefix to the Global Mirror name created by the job.

Note: Global Mirror labels must contain only alphanumeric characters and underscores.

Destination storage limit in GB / Destination volumes limit

Specify quotas for storage usage and the number of volume created on the destination for all jobs utilizing the SLA Policy.

Configure IBM Spectrum Accelerate SLA Policies

To add a VM Replication sub-policy to an IBM Spectrum Accelerate SLA Policies:

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.
Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.
2. Click **Add VM Copy** .
3. In the Destination pane select an IBM host destination from the list of available resources as the VM Replication destination, along with an associated storage pool. If no storage pool is selected, the storage pool with the largest amount of space available is chosen by default. To select the original target destination, select **Use Original**.
4. In the Options pane set the VM Replication sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Golden Snapshot

Enable to create a golden snapshot on a thin provisioned pool, which cannot be deleted through an automated process .

Note that when space is limited, IBM Spectrum Accelerate snapshots are deleted after associated jobs complete. In a thin provisioned pool, the golden snapshot option ensures the snapshot will not be deleted. This option is not compatible with thick provisioned pools, and may lead to the loss of data if the SLA Policy snapshot is expired prematurely by the storage system to reclaim space.

Target Volume Prefix Label

Enter an optional label to identify the target volume. This label is added as a prefix to the volume name created by the job.

Note: Volume prefix labels must contain only alphanumeric characters and underscores. Labels cannot begin with numeric characters.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Note: Snapshot labels must contain only alphanumeric characters and underscores.

Name

Enter an optional label to replace the default snapshot sub-policy label displayed in IBM Spectrum Copy Data Management. The default initial label is VM Copy0.

Protocol

If more than one storage protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

Full Copy Method

Select the full copy method. Available full copy methods include Clone or VADP-based VM Replication.

Destination storage limit in GB / Destination volumes limit

Specify quotas for storage usage and the number of volume created on the destination for all jobs utilizing the SLA Policy.

To add a snapshot sub-policy to an IBM Spectrum Accelerate SLA Policy:

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add Snapshot** . In the Options pane, set the snapshot sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Golden Snapshot

Enable to create a golden snapshot on a thin provisioned pool, which cannot be deleted through an automated process .

Note that when space is limited, IBM Spectrum Accelerate snapshots are deleted after associated jobs complete. In a thin provisioned pool, the golden snapshot option ensures the snapshot will not be deleted. This option is not compatible with thick provisioned pools, and may lead to the loss of data if the SLA Policy snapshot is expired prematurely by the storage system to reclaim space.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Note: Snapshot labels must contain only alphanumeric characters and underscores.

Name

Enter an optional name to replace the default snapshot sub-policy name displayed in IBM Spectrum Copy Data Management. The default initial name is Snapshot0.

To add a Replication sub-policy to an IBM Spectrum Accelerate SLA Policy:

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add Replication**.
3. In the Replication Destination pane select an IBM Spectrum Accelerate host destination from the list of available resources as the Replication destination, along with an associated storage pool. If no storage pool is selected, the storage pool with the largest amount of space available is chosen by default.
4. In the Options pane set the Replication sub-policy options.

Keep Source Snapshots / Keep Destination Snapshots

A replication sub-policy provides snapshots to both a Source, or Primary location, and a Destination, or Replication location. After a certain number of snapshot instances are created for a resource, older instances are purged from the Source and Destination. In the **Keep Source Snapshots** and **Keep Destination Snapshots** fields, enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots (maximum)** field.

Golden Snapshot

Enable to create a golden snapshot on a thin provisioned pool, which cannot be deleted through an automated process .

Note that when space is limited, IBM Spectrum Accelerate snapshots are deleted after associated jobs complete. In a thin provisioned pool, the golden snapshot option ensures the snapshot will not be deleted. This option is not compatible with thick provisioned pools, and may lead to the loss of data if the SLA Policy snapshot is expired prematurely by the storage system to reclaim space.

Name

Enter an optional name to replace the default Replication sub-policy name displayed in IBM Spectrum Copy Data Management. The default initial name is Replication0.

Keep Source Volume name for target volume

Enable to retain the source volume name for copy data generated by IBM Spectrum Copy Data Management.

Target Volume Prefix Label

Enter an optional label to identify the target volume. This label is added as a prefix to the volume name created by the job.

Note: Volume prefix labels must contain only alphanumeric characters and underscores. Labels cannot begin with numeric characters.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Note: Snapshot labels must contain only alphanumeric characters and underscores.

Synchronization

Specify the time in minutes in which the change volumes will be synchronized with a consistent copy of the data. If a copy does not complete in the cycle period, the next cycle period will not start until the copy is complete. Synchronization can also be initiated manually.

Destination storage limit in GB / Destination volumes limit

Specify quotas for storage usage and the number of volume created on the destination for all jobs utilizing the SLA Policy.

Configure NetApp ONTAP SLA Policies

To add a snapshot sub-policy to a NetApp ONTAP SLA Policy:

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add Snapshot** . In the Options pane, set the snapshot sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Disable system snapshot policy

Disables all the system snapshot jobs on the storage volumes.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Note: Snapshot labels must contain only alphanumeric characters and underscores.

Name

Enter an optional name to replace the default snapshot sub-policy name displayed in IBM Spectrum Copy Data Management. The default initial name is Snapshot0.

To add a VM Replication sub-policy to a NetApp ONTAP SLA Policy:

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add VM Copy** .
3. In the VM Replication Destination pane select an SVM from the list of available resources as the VM Replication destination, along with an associated aggregate. If no aggregate is selected, the aggregate with the largest amount of space available is chosen by default.
4. In the Options pane set the VM Replication sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Disable system snapshot policy

Disables all the system snapshot jobs on the storage volumes.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Note: Snapshot labels must contain only alphanumeric characters and underscores.

Name

Enter an optional label to replace the default snapshot sub-policy label displayed in IBM Spectrum Copy Data Management. The default initial label is VM Copy0.

Storage Efficiency (Deduplication)

Enable or disable storage efficiency. Storage efficiency uses data deduplication to store the maximum amount of data while consuming less space.

Destination Datastore Type

Set the destination datastore type. Available datastore types include NFS and VMFS.

NFS - A NetApp volume will be created for NFS access and the target datastore will be created on that NFS share.

VMFS - A NetApp volume will be created and a LUN will be created on the volume. The volume will be mapped to the ESX, and the LUN will be formatted for VMFS. A VMFS datastore will be created on the LUN.

To add a mirror sub-policy to a NetApp ONTAP SLA Policy:

1. Select a snapshot, VM Replication, vault or mirror from the workflow pane and click **Add Mirror** .
2. In the Mirror Destination pane select a storage controller or SVM from the list of available resources as the mirror destination, along with an associated aggregate. If no aggregate is selected, the aggregate with the largest amount of space available is chosen by default.
3. In the Options pane set the **Mirror**  sub-policy options.

Name

Enter an optional name to replace the default mirror sub-policy name displayed in IBM Spectrum Copy Data Management. The default initial name is Mirror0.

Keep Source Volume name for target volume

Enable to retain the source volume name for copy data generated by IBM Spectrum Copy Data Management.

Volume Prefix Label

Enter an optional label to identify the volume. This label is added as a prefix to the volume name created by the job and cannot be edited after the job is submitted.

Note: Volume prefix labels must contain only alphanumeric characters and underscores. Labels cannot begin with numeric characters.

Storage Efficiency (Deduplication)

Enable or disable snapshot storage efficiency. Storage efficiency uses data deduplication to store the maximum amount of data while consuming less space.

Throttle

Set the transfer throughput in KBs per second between the source and the destination, which controls the number of parallel transfers that can take place.

Destination storage limit in GB / Destination volumes limit

Specify quotas for storage usage and the number of volume created on the destination for all jobs utilizing the SLA Policy.

To add a vault sub-policy to a NetApp ONTAP SLA Policy:

1. Select a snapshot, vault or mirror from the workflow pane and click **Add Vault** .
2. In the Vault Destination pane select a storage controller or SVM from the list of available resources as the vault destination, along with an associated aggregate. If no aggregate is selected, the aggregate with the largest amount of space available is chosen by default.
3. In the Options pane set the **Vault**  sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Name

Enter an optional name to replace the default vault sub-policy name displayed in IBM Spectrum Copy Data Management. The default initial name is Vault0.

Keep Source Volume name for target volume

Enable to retain the source volume name for copy data generated by IBM Spectrum Copy Data Management.

Volume Prefix Label

Enter an optional label to identify the volume. This label is added as a prefix to the volume name created by the job and cannot be edited after the job is submitted.

Note: Volume prefix labels must contain only alphanumeric characters and underscores. Labels cannot begin with numeric characters.

Storage Efficiency (Deduplication)

Enable or disable snapshot storage efficiency. Storage efficiency uses data deduplication to store the maximum amount of data while consuming less space.

Throttle

Set the transfer throughput in KBs per second between the source and the destination, which controls the number of parallel transfers that can take place.

Destination storage limit in GB / Destination volumes limit

Specify quotas for storage usage and the number of volume created on the destination for all jobs utilizing the SLA Policy.

Configure HPE Nimble Storage SLA Policies

To add a snapshot sub-policy to a HPE Nimble Storage SLA Policy:

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add Snapshot** . In the Options pane, set the snapshot sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Note: Snapshot labels must contain only alphanumeric characters and underscores.

Name

Enter an optional name to replace the default snapshot sub-policy name displayed in IBM Spectrum Copy Data Management. The default initial name is Snapshot0.

Configure Pure Storage FlashArray SLA Policies

To add a VM Replication sub-policy to a Pure Storage FlashArray SLA Policy:

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add VM Copy** .
3. In the VM Replication Destination pane select a Pure Storage FlashArray host destination from the list of available resources as the VM Replication destination.
4. In the Options pane set the VM Replication sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Target Volume Prefix Label

Enter an optional label to identify the target volume. This label is added as a prefix to the volume name created by the job.

Note: Volume prefix labels must contain only alphanumeric characters and underscores. Labels cannot begin with numeric characters.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Note: Snapshot labels must contain only alphanumeric characters and underscores.

Name

Enter an optional label to replace the default snapshot sub-policy label displayed in IBM Spectrum Copy Data Management. The default initial label is VM Copy0

Protocol

If more than one storage protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

Full Copy Method

Select the full copy method. Available methods include Clone or VADP-based VM Replication.

To add a snapshot sub-policy to a Pure Storage FlashArray SLA Policy:

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add Snapshot** . In the Options pane, set the snapshot sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Note: Snapshot labels must contain only alphanumeric characters and underscores.

Name

Enter an optional name to replace the default snapshot sub-policy name displayed in IBM Spectrum Copy Data Management. The default initial name is Snapshot0.

3. If your Pure Storage FlashArray supports CloudSnap functionality, you can add snapshot offload which will create an offload copy on cloud storage. This option is only available after a snapshot sub-policy is added. Click **Add Snapshot Offload**  or right-click on the snapshot sub-policy and click **Add Snapshot Offload** . In the Options pane, set the snapshot offload sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Best Practice: The ECX user interface will indicate that a Pure Storage FlashArray CloudSnap offload job has completed even though the transfer is still occurring in the background which is dependent on network speeds. Consider setting an age as the retention for offload copies when using the Pure Storage FlashArray CloudSnap functionality. Doing so will ensure that a sufficient amount of time has passed for data to be transferred to the S3 storage target before it has to be condensed out from a backup. This is particularly important if several offload jobs are run in quick succession.

Offload Target

Select **Cloud** from the menu as the offload target.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Note: Snapshot labels must contain only alphanumeric characters and underscores.

Name

Enter an optional name to replace the default snapshot sub-policy name displayed in IBM Spectrum Copy Data Management. The default initial name is Snapshotoffload0.

To add a Replication sub-policy to a Pure Storage FlashArray SLA Policy:

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add Replication** .
3. In the Replication Destination pane select a Pure Storage FlashArray host destination from the list of available resources as the Replication destination.
4. In the Options pane set the Replication sub-policy options.

Keep Source Snapshots / Keep Destination Snapshots

A Pure Storage replication sub-policy provides snapshots to both a Source, or Primary location, and a Destination, or Replication location. After a certain number of snapshot instances are created for a resource, older instances are purged from the Source and Destination. In the **Keep Source Snapshots** and **Keep Destination Snapshots** fields, enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots (maximum)** field.

Name

Enter an optional name to replace the default Replication sub-policy name displayed in IBM Spectrum Copy Data Management. The default initial name is Replication0.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Note: Snapshot labels must contain only alphanumeric characters and underscores.

5. When you are satisfied that the SLA Policy-specific information is correct, click **Finish**. The SLA Policy appears on the All SLA Policies pane and can be applied to new and existing Backup job definitions.

NEXT STEPS:

- Create a Backup job definition that utilizes an SLA Policy.

RELATED TOPICS:

- [Create a Backup Job Definition - DellEMC Unity](#) on page **220**
- [Create a Backup Job Definition - IBM Spectrum Virtualize](#) on page **227**
- [Create a Backup Job Definition - NetApp ONTAP](#) on page **231**
- [Create a Backup Job Definition - Pure Storage FlashArray](#) on page **238**
- [Create a Backup Job Definition - VMware](#) on page **242**

Configure Scripts

Prescripts and postscripts are scripts that can be run before or after Backup and Restore jobs run, both at a job-level and before or after snapshots are captured. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

Scripts can be created locally, uploaded to your environment through the Scripts pane, then applied to job definitions. In a Windows environment, if your application supports VSS, the Backup job triggers the VSS application quiesce logic if the **Make these VMs application/file system consistent** option is enabled when creating the VMware Backup job. However, for applications that don't support VSS, or on Linux virtual machines, pre and post snapshot scripts can be used to quiesce your application for the snapshot backup.

Note: If adding a script to a Windows-based File System job definition, the user running the script must have the "Log on as a service" right enabled, which is required for running prescripts and postscripts. For more information about the "Log on as a service" right, see <https://technet.microsoft.com/en-us/library/cc794944.aspx>.

Upload a Script

Supported scripts include shell scripts for Linux-based virtual machines and Batch and PowerShell scripts for Windows-based virtual machines. Scripts must be created using the operating system's associated file format.

1. Click the **Configure**  tab. On the Views pane, select **Scripts**  .
2. Click **Upload**  . The Upload Script dialog opens.
3. In the Script field, browse for a local script to upload, then click **Open**.
4. Enter an optional comment, then click **OK**. The script appears on the **Scripts** pane and can be applied to supported jobs.

Replace a Script

Upload a revised version of a script.

1. Click the **Configure**  tab. On the Views pane, select **Scripts**  .
2. In the Scripts pane, select the script to replace by clicking in the row containing the script name.
3. Click **Replace**  . The Update Script dialog opens.
4. In the Script field, browse for a local updated script to upload, then click **Open**.
5. Enter an optional comment, then click **OK**. The revised script appears on the **Scripts** pane and can be applied to supported jobs.

Delete a Script

Delete a script when it becomes obsolete. Removing a script from an associated job definition allows you to delete the script immediately. Deleting the job definition while a script is still assigned to a job definition requires that you run the Maintenance job before deleting the script.

1. Click the **Configure**  tab. On the Views pane, select **Scripts**  .
2. In the Scripts pane, select the script to delete by clicking in the row containing the script name.
3. Click **Delete** . A confirmation dialog box displays.
4. Confirm deletion. The script is deleted.

RELATED TOPICS:

- [Job Definition Overview](#) on page **166**
- [Using State and Status Arguments in Postscripts](#) on page **334**
- [Return Code Reference](#) on page **553**

Schedules

The topics in the following section cover creating, editing, and deleting schedules.

Create a Schedule

A schedule is a set of rules for triggering a job. Create a schedule to apply to one or more jobs. Once applied, the job sessions are run as defined by the parameters of the schedule.

Best Practice: Overlapping schedules may slow down your network. Decrease the strain on your network by configuring multiple schedules to run at different times or days of the week.

BEFORE YOU BEGIN:

- View the schedules that are already set up to determine if you have one that suits your requirements. See [Edit a Schedule](#) on page 157.

To create a schedule:

1. Click the **Configure**  tab. On the Views pane, select **Schedules** .
2. In the Schedules pane, click **New** . The Properties pane displays.
3. Revise fields on the Properties pane:
 - a. In **Name**, enter a descriptive schedule name. By default, your schedule parameters are added to the schedule description.
 - b. In **Trigger**, select the frequency for the job session to run:
 - Once** to schedule a job session to run once. In **Trigger date**, select the day of the week for the job session to run. In **Time of day**, select a starting time.
 - Hourly** to schedule a job session to run hourly. In **Interval**, select the number of hours between job sessions. In **Time of day** and **Starts**, select a starting time and date. If applicable, enter an expiration date in **Expires**.
 - Daily** to schedule a job session to run daily or every few days. In **Interval**, select the number of days between job sessions. In **Time of day** and **Starts**, select a starting time and date. If applicable, enter an expiration date in **Expires**.
 - Weekly** to schedule a job session to run weekly or every few weeks. In **Day of week**, select the day of the week for the job session to run during the week. In **Time of day** and **Starts**, select a starting time and date. If applicable, enter an expiration date in **Expires**.
 - Monthly** to schedule a job session to run monthly or every few months. In **Day of month**, select the day or days of the month for the job session to run during the month. In **Time of day** and **Starts**, select a starting time and date. If applicable, enter an expiration date in **Expires**.
4. Click **Save** . The schedule displays in the Schedules pane.

NEXT STEPS:

- Assign the schedule to a new or existing job definition. See [Jobs Overview](#) on page **161** and [Edit a Job Definition](#) on page **342**.

RELATED TOPICS:

- [Edit a Schedule](#) on page **157**
- [Delete a Schedule](#) on page **159**

Edit a Schedule

Revise a schedule to change the timetable for running a job session. Because a single schedule can be applied to multiple jobs, all jobs associated with the schedule you are editing are impacted.

Best Practice: Overlapping schedules may slow down your network. Decrease the strain on your network by configuring multiple schedules to run at different times or days of the week.

BEFORE YOU BEGIN:

- Review the properties of your current schedules. See [Create a Schedule](#) on page 155.

To edit a schedule:

1. Click the **Configure**  tab. On the Views pane, select **Schedules** .
2. Select the schedule to edit by clicking in the row containing the schedule name.
3. Revise fields on the Properties pane:
 - a. In **Name**, enter a descriptive schedule name. By default, your schedule parameters are added to the schedule description.
 - b. In **Name**, enter a descriptive schedule name. By default, your schedule parameters are added to the schedule description.
 - Once** to schedule a job session to run once. In **Trigger date**, select the day of the week for the job session to run. In **Time of day**, select a starting time.
 - Hourly** to schedule a job session to run hourly. In **Interval**, select the number of hours between job sessions. In **Time of day** and **Starts**, select a starting time and date. If applicable, enter an expiration date in **Expires**.
 - Daily** to schedule a job session to run daily or every few days. In **Interval**, select the number of days between job sessions. In **Time of day** and **Starts**, select a starting time and date. If applicable, enter an expiration date in **Expires**.
 - Weekly** to schedule a job session to run weekly or every few weeks. In **Day of week**, select the day of the week for the job session to run during the week. In **Time of day** and **Starts**, select a starting time and date. If applicable, enter an expiration date in **Expires**.
 - Monthly** to schedule a job session to run monthly or every few months. In **Day of month**, select the day or days of the month for the job session to run during the month. In **Time of day** and **Starts**, select a starting time and date. If applicable, enter an expiration date in **Expires**.
4. Click **Save** . The revisions are applied to the schedule.

NEXT STEPS:

- Assign the edited schedule to a new or existing job definition. See [Jobs Overview](#) on page **161** and [Edit a Job Definition](#) on page **342**.

RELATED TOPICS:

- [Create a Schedule](#) on page **155**
- [Delete a Schedule](#) on page **159**

Delete a Schedule

Delete a schedule from the application if it is not used to trigger jobs.

To delete a schedule:

1. Click the **Configure**  tab. On the Views pane, select **Schedules** .
2. Select the schedule to delete by clicking in the row containing the schedule name.
3. Click **Delete** . A confirmation dialog box opens.
4. Confirm deletion. The schedule is deleted.

RELATED TOPICS:

- [Create a Schedule](#) on page **155**
- [Edit a Schedule](#) on page **157**

Jobs

The topics in the following section cover defining, editing, and deleting job definitions, as well as descriptions of job types.

Jobs Overview

From the **Jobs**  tab you can create and edit job definitions. A job definition is a user-defined set of tasks and rules. Once a job definition is added to IBM Spectrum Copy Data Management, it can be combined with a schedule or trigger to create a job. There are several job types including Inventory, Backup, Restore, Reports, and Scripts.

You can also start, monitor, stop, and resume job sessions from the **Jobs**  tab. From this pane, you can also view all scheduled and unscheduled job sessions, and start jobs before scheduled run times.

Note: If the IBM Spectrum Copy Data Management virtual appliance is shut down while jobs are in progress, the jobs will automatically restart once the virtual appliance is back online. By default, jobs that began running 30 minutes before the appliance restarted will automatically restart.

In addition, you can view the job details. Select a job to view the current job status, the job schedule, and control the activity for the selected job.

Select a job and click **View Last Run Session**  to view the duration and completion status of the most recent run of a selected job.

- [Job Definition Overview](#) on page **166**
- [Monitor a Job Session](#) on page **164**
- [Start, Pause, and Hold a Job Session](#) on page **162**

Start, Pause, and Hold a Job Session

From the **Jobs**  tab you can run a job session on demand, pause or cancel a running job, and hold all future scheduled instances of a job from running until you are ready for the job to proceed.

If the IBM Spectrum Copy Data Management virtual appliance is shut down while jobs are in progress, the jobs will automatically restart once the virtual appliance is back online. By default, jobs that began running 30 minutes before the appliance restarted will automatically restart.

BEFORE YOU BEGIN:

- Create job definitions and job schedules. See [Jobs Overview](#) on page 161.

To start a job session:

1. Click the **Jobs**  tab.
2. Select the job to run by clicking in the row containing the job name.
3. Click **Start** , or right-click the job name and select **Start** . A confirmation dialog box opens.

Note: If a job session has multiple run options, such as running a job session in Test, Recovery, or Clone mode, you will be prompted to select a job session type.

4. Click **Yes**. The job session runs.
5. In the **Activity** pane, click the job name to view the job session details, including the job session's start date and time, duration, description, status through a progress bar, and associated messages.

To pause and resume a job session:

1. Click the **Jobs**  tab.
2. Select a running job to stop by clicking in the row containing the job name.
3. From the **More Actions** drop-down menu, select **Pause** , or right-click the job name and select **Pause** . A confirmation dialog box opens.
4. Click **Yes**. The job session pauses.
5. Select **Resume**  from the **More Actions** drop-down menu, or right-click the job name and select **Resume**  to resume the job session.

To hold and release a job's schedule:

1. Click the **Jobs**  tab.
2. Select the job to suspend by clicking in the row containing the job name.

3. From the **More Actions** drop-down menu, select **Hold Schedule**, or right-click the job name and select **Hold Schedule** to hold the job. The job session status changes to **Held**, and all future scheduled instances of the job will not run until released.
4. From the **More Actions** drop-down menu, select **Release Schedule**, or right-click the job name and select **Release Schedule** to release the job

To cancel a job session:

1. Click the **Jobs**  tab.
2. Select the job session to cancel by clicking in the row containing the job name.
3. From the **More Actions** drop-down menu, select **Cancel**  to cancel the job. The job session status changes to **Canceled**.

NEXT STEPS:

- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- Once the job session completes, review cataloged data through the Search and Report tabs. See [Browse Inventory](#) on page **357** and [Report Overview](#) on page **360**.

RELATED TOPICS:

- [Create a Schedule](#) on page **155**

Monitor a Job Session

You can view the details of a job session that is currently running or one that has finished. Use the **Jobs**  tab to view the status of a job session including start time, end time, and job name.

BEFORE YOU BEGIN:

- Start a job session. If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.

To monitor a running job session:

1. Click the **Jobs**  tab. The All Jobs folder displays defined jobs that are currently running or idle, and provides information about their most recent session, their last runtime, last run duration, and last run status. To view a specific job type, select the associated job type from the folder structure.
2. View the status of a job in the status column. Currently running job sessions are represented by an active icon. Once a job session finishes, one of the following icons appears in the status column:
 -  **Completed** - Indicates the job session completed successfully. All tasks associated with the job session were completed.
 -  **Partial** - Indicates the job session completed, but one or more tasks failed or were skipped.
 -  **Failed** - Indicates the job session did not successfully complete due to mixed task statuses.
 -  **Aborted** - Indicates the job session did not successfully complete due to a reset, reboot, or shutdown of the virtual appliance server.
 -  **Held** - Indicates the job has been paused through the Halt feature in the Actions menu.
 -  **Idle** - Indicates the job session is idle.
 -  **Skipped** - Indicates that a volume was not cataloged. See the Task tab for more information about skipped jobs.
 -  **Stopped** - Indicates the job was stopped using the Stop button.

To filter the list of jobs based on type or status:

1. Click the **Jobs**  tab.
2. Click the drop-down arrow in the header of the **Status** or **Last Run Status** columns.
3. Select **Filters**, then choose a filter criteria.

To view information about specific job sessions:

1. Click the **Jobs**  tab and select a job by clicking in the row containing the job name.
2. The Activity / History pane displays. If the job is currently running, its status displays in the **Activity** tab. To review the history of previous runs of the job, click the **History** tab, then click the session for details. The following panes open:

Tasks

Displays a task-by-task view of the job session, including start and end times, duration, and status. It also displays details of the underlying tasks that take place during the job session, including the task's type, duration, and status.

Details

Displays an overview of the job definition, including the job name, sources, options, and notification settings.

Log

Displays the job log, which can be used for troubleshooting purposes.

RELATED TOPICS:

- [Start, Pause, and Hold a Job Session](#) on page **162**
- [Collect Logs For Troubleshooting](#) on page **518**

Job Definition Overview

From the **Jobs**  tab you can create and edit job definitions. A job definition is a user-defined set of tasks and rules. Once a job definition is added to IBM Spectrum Copy Data Management, it can be combined with a schedule or trigger to create a job. There are several job types including Inventory, Backup, Restore, Reports, and Scripts.

A schedule is a set of rules for triggering a job definition. Create a schedule to apply to one or more job. Once a schedule is applied, the job sessions are run as defined by the parameters of the schedule.

When a job is run, the job session status is monitored and its status can be watched real-time in the job monitor. Job sessions are run on demand or by trigger.

Best Practice: Create a schedule before creating a job definition so that you can easily add the schedule to the job definition.

Job Types

Inventory jobs

Inventory jobs interrogate storage systems to gather and record metadata about high-level objects and files. You can select one or more providers of the same type in a single job definition for cataloging.

Backup and Restore jobs

IBM Spectrum Copy Data Management utilizes automated Copy Data Management workflows for replicating and intelligently reusing snapshots, vaults, and mirrors. Backup and Restore jobs offer control over testing and cloning use cases, instant recovery, and full disaster recovery. Through Backup and Restore jobs, you can:

- Copy data from a variety of storage providers to multiple locations.
- Reuse and recover resources from snapshots, vaults, mirrors, and other copies and replicas.
- Support use cases for automated data protection, recovery, DevOps, Dev/Test, data and database validation with data masking, through the use of automated Instant Disk Restore, Instant VM Restore, volume, and file restore functionalities.

Report jobs

A Report job is a System job that summarizes information about cataloged providers and the data and other resources that reside on them.

Script jobs

A Script job defines a set of commands to run on the IBM Spectrum Copy Data Management appliance. Use the script job to add functionality to IBM Spectrum Copy Data Management. A script can consist of one or many commands, such as a shell script.

Maintenance job

The Maintenance job removes resources and associated objects created by IBM Spectrum Copy Data Management when a job in a pending state is deleted. The cleanup procedure reclaims space on your storage devices, cleans up your IBM Spectrum Copy Data Management catalog, and removes related snapshots.

RELATED TOPICS:

- [Configure SLA Policies](#) on page **133**
- [Create a Schedule](#) on page **155**
- [Monitor a Job Session](#) on page **164**

Inventory Jobs

- [Create an Inventory Job Definition - Database](#) on page **170**
- [Create an Inventory Job Definition - File System](#) on page **173**
- [Create an Inventory Job Definition - DellEMC Unity](#) on page **175**
- [Create an Inventory Job Definition - IBM Spectrum Accelerate](#) on page **177**
- [Create an Inventory Job Definition - IBM Spectrum Protect Snapshot](#) on page **179**
- [Create an Inventory Job Definition - IBM Spectrum Virtualize](#) on page **181**
- [Create an Inventory Job Definition - NetApp ONTAP Storage](#) on page **183**
- [Create an Inventory Job Definition - NetApp ONTAP File](#) on page **186**
- [Create an Inventory Job Definition - HPE Nimble Storage](#) on page **189**
- [Create an Inventory Job Definition - Pure Storage FlashArray](#) on page **192**
- [Create an Inventory Job Definition - VMware](#) on page **194**

Backup Jobs

- [Create a Backup Job Definition - InterSystems Caché](#) on page **197**
- [Create a Backup Job Definition - SAP HANA](#) on page **200**
- [Create a Backup Job Definition - Oracle](#) on page **204**
- [Create a Backup Job Definition - SQL](#) on page **211**
- [Create a Backup Job Definition - File System](#) on page **216**
- [Create a Backup Job Definition - DellEMC Unity](#) on page **220**
- [Create a Backup Job Definition - IBM Spectrum Accelerate](#) on page **223**
- [Create a Backup Job Definition - IBM Spectrum Virtualize](#) on page **227**
- [Create a Backup Job Definition - NetApp ONTAP](#) on page **231**
- [Create a Backup Job Definition - HPE Nimble Storage](#) on page **234**
- [Create a Backup Job Definition - Pure Storage FlashArray](#) on page **238**
- [Create a Backup Job Definition - VMware](#) on page **242**

Restore Jobs

- [Create a Restore Job Definition - InterSystems Caché](#) on page **252**
- [Create a Restore Job Definition - SAP HANA](#) on page **257**
- [Create a Restore Job Definition - Oracle](#) on page **264**
- [Create a Restore Job Definition - SQL](#) on page **270**
- [Create a Restore Job Definition - File System](#) on page **277**
- [Create a Restore Job Definition - DellEMC Unity](#) on page **281**
- [Create a Restore Job Definition - IBM Spectrum Accelerate](#) on page **287**
- [Create a Restore Job Definition - IBM Spectrum Virtualize](#) on page **293**
- [Create a Restore Job Definition - NetApp ONTAP](#) on page **299**
- [Create a Restore Job Definition - HPE Nimble Storage](#) on page **306**
- [Create a Restore Job Definition - Pure Storage FlashArray](#) on page **312**
- [Create a Restore Job Definition - VMware](#) on page **317**

System Jobs

- [Maintenance Job](#) on page **336**
- [Create a Report Job Definition](#) on page **337**
- [Create a Script Job Definition](#) on page **339**

Inventory Jobs

The topics in the following section cover Inventory job definitions.

Create an Inventory Job Definition - Database

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

A Database Inventory job provides the framework to catalog and collect information about your application database servers.

Run a Database Inventory job to analyze your application servers in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer sprawl, storage overutilization, and other storage inefficiencies.

Note: To successfully catalog virtual providers, you must also register associated vCenter providers.

BEFORE YOU BEGIN:

- At least one application provider must be associated with a Database Inventory job definition, such as an Oracle resource. Before defining an Inventory job, add application providers. See [Register a Provider](#) on page 75.
- Credentials are required for cataloging application servers. See [Identities Overview](#) on page 127.
- Review Oracle and SQL requirements. See [Oracle Requirements](#) on page 41, [Microsoft SQL Server Requirements](#) on page 57, [SAP HANA Requirements](#) on page 51, [InterSystems Caché Requirements](#) on page 34.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- If an Oracle Inventory job runs at the same time or short period after an Oracle Backup job runs, copy errors may occur due to temporary mounts that are created during the Backup job. As a best practice, schedule Oracle Inventory jobs so that they do not overlap with Oracle Backup jobs.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create a Database Inventory job definition:

1. Click the **Jobs**  tab. Expand the **Database** folder, then select an application database type to catalog. Available application database types include **InterSystems Caché**, **Oracle**, **SAP HANA**, and **Microsoft SQL**.
2. Click **New** , then select **Inventory**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.

- From the list of available sites, select one or more providers containing application data that you wish to catalog.
- To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
- To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

Number of catalog instances to keep

After a certain number of job runs for a given job, older objects for that job are purged from the Inventory.

Enter the number of job runs for which high-level objects are to be retained.

- Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

- Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

- When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

Note: If you selected the **Start job now** option, the job runs.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**

Create an Inventory Job Definition - File System

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

A File System Inventory job provides the framework to catalog and collect information about your physical Windows, Linux and AIX file systems. You can select one or more file system providers of the same type in a single job definition for cataloging.

Run a File System Inventory job to analyze your file systems in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer sprawl, storage overutilization, and other storage inefficiencies.

BEFORE YOU BEGIN:

- Review File System requirements. See [File System Requirements](#) on page 28.
- At least one file system provider must be associated with a File System Inventory job definition. Before defining an Inventory job, add file system providers. See [Register a Provider](#) on page 75.
- Credentials are required for cataloging application servers. See [Identities Overview](#) on page 127.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create a File System Inventory job definition:

1. Click the **Jobs**  tab. Expand the **File System** folder, then select **File System**.
2. Click **New** , then select **Inventory**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more file systems that you wish to catalog.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

Number of catalog instances to keep

After a certain number of job runs for a given job, older objects for that job are purged from the Inventory. Enter the number of job runs for which high-level objects are to be retained.

7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

8. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

9. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

Note: If you selected the **Start job now** option, the job runs.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page 162.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page 164.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Edit a Job Definition](#) on page 342
- [Delete a Job Definition](#) on page 343
- [Create a Schedule](#) on page 155

Create an Inventory Job Definition - DellEMC Unity

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

A DellEMC Unity Inventory job provides the framework to catalog and collect information about high-level DellEMC Unity objects on your DellEMC Unity storage systems. You can select one or more DellEMC Unity providers in a single job definition for cataloging.

Run a DellEMC Unity Inventory job to analyze your DellEMC Unity environment in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer sprawl, storage overutilization, and other storage inefficiencies.

BEFORE YOU BEGIN:

- At least one DellEMC Unity storage provider must be associated with a DellEMC Unity Inventory job definition. Before defining an Inventory job, add DellEMC Unity storage providers. See [Register a Provider](#) on page 75.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create a DellEMC Unity Inventory job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **DellEMC Unity**.
2. Click **New** , then select **Inventory**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more providers to catalog.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

Number of catalog instances to keep

After a certain number of job runs for a given job, older DellEMC Unity objects for that job are purged from the Inventory. Enter the number of job runs for which high-level DellEMC Unity objects are to be retained.

7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

8. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

9. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

Note: If you selected the **Start job now** option, the job runs.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**

Create an Inventory Job Definition - IBM Spectrum Accelerate

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

An IBM Spectrum Accelerate Inventory job provides the framework to catalog and collect information about high-level objects on your IBM Spectrum Accelerate storage systems. You can select one or more IBM Spectrum Accelerate providers in a single job definition for cataloging.

Run an IBM Spectrum Accelerate Inventory job to analyze your IBM Spectrum Accelerate environment in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer sprawl, storage overutilization, and other storage inefficiencies.

BEFORE YOU BEGIN:

- At least one IBM Spectrum Accelerate storage provider must be associated with an IBM Spectrum Accelerate Inventory job definition. Before defining an Inventory job, add IBM Spectrum Accelerate storage providers. See [Register a Provider](#) on page 75.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create an IBM Spectrum Accelerate Inventory job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **IBM Spectrum Accelerate**.
2. Click **New** , then select **Inventory**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more providers to catalog.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

Number of catalog instances to keep

After a certain number of job runs for a given job, older IBM Spectrum Accelerate objects for that job are purged from the Inventory. Enter the number of job runs for which high-level IBM Spectrum Accelerate objects are to be retained.

- Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

- Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

- When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

Note: If you selected the **Start job now** option, the job runs.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**

Create an Inventory Job Definition - IBM Spectrum Protect Snapshot

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

An IBM Spectrum Protect Snapshot Inventory job provides the framework to catalog and collect information about high-level IBM objects. You can select one or more IBM Spectrum Protect providers in a single job definition for cataloging.

Run an IBM Spectrum Protect Snapshot Inventory job to analyze your IBM environment in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer sprawl, storage overutilization, and other storage inefficiencies.

Note: To successfully catalog an IBM Spectrum Protect Snapshot provider, you must also register associated vCenters and IBM Storage providers. IBM providers utilize port 22 for communication with IBM Spectrum Copy Data Management.

BEFORE YOU BEGIN:

- At least one IBM Spectrum Protect Snapshot provider must be associated with an IBM Spectrum Protect Snapshot Inventory job definition. Before defining an Inventory job, add an IBM Spectrum Protect Snapshot provider. See [Register a Provider](#) on page 75.
- Register vCenters and IBM storage providers associated with the IBM Spectrum Protect Snapshot provider.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create an IBM Spectrum Protect Snapshot Inventory job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **IBM Spectrum Protect Snapshot**.
2. Click **New** , then select **Inventory**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more providers to catalog.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.

6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.
7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

8. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.
Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.
9. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

Note: If you selected the **Start job now** option, the job runs.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**

Create an Inventory Job Definition - IBM Spectrum Virtualize

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

An IBM Spectrum Virtualize Inventory job provides the framework to catalog and collect information about high-level IBM objects on your IBM storage systems. You can select one or more IBM providers in a single job definition for cataloging.

Run an IBM Spectrum Virtualize Inventory job to analyze your IBM environment in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer sprawl, storage overutilization, and other storage inefficiencies.

BEFORE YOU BEGIN:

- At least one IBM storage provider must be associated with an IBM Spectrum Virtualize Inventory job definition. Before defining an Inventory job, add IBM storage providers. See [Register a Provider](#) on page 75.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- IBM providers utilize port 22 for communication with IBM Spectrum Copy Data Management.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create an IBM Spectrum Virtualize Inventory job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **IBM Spectrum Virtualize**.
2. Click **New** , then select **Inventory**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more providers to catalog.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

Number of catalog instances to keep

After a certain number of job runs for a given job, older IBM objects for that job are purged from the Inventory. Enter the number of job runs for which high-level IBM objects are to be retained.

7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

8. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

9. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

Note: If you selected the **Start job now** option, the job runs.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page 162.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page 164.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Edit a Job Definition](#) on page 342
- [Delete a Job Definition](#) on page 343
- [Create a Schedule](#) on page 155

Create an Inventory Job Definition - NetApp ONTAP Storage

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

A NetApp ONTAP Storage Inventory job provides the framework to catalog and collect information about high-level NetApp ONTAP objects. You can select one or more NetApp ONTAP cluster or non-cluster providers in a single job definition for cataloging. NetApp ONTAP 7-Mode and Cluster-Mode are both supported.

Run a NetApp ONTAP Storage Inventory job to analyze your NetApp ONTAP environment in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer snapshot sprawl, storage overutilization, and other storage inefficiencies.

Note: For cataloging low-level NetApp ONTAP objects, create a NetApp ONTAP File Inventory job definition.

BEFORE YOU BEGIN:

- At least one NetApp ONTAP provider must be associated with a NetApp ONTAP Storage Inventory job definition. Before defining an Inventory job definition, add NetApp ONTAP providers. See [Register a Provider](#) on page 75.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create a NetApp ONTAP Storage Inventory job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **NetApp ONTAP**.
2. Click **New** , then select **Inventory NetApp ONTAP Storage**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more providers to catalog.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

Connection timeout (secs)

To run a catalog job, the application needs to connect with the resource. If there is no response within a certain time limit, it times out and the job session fails. Enter the number of seconds to wait before timing out.

Number of catalog instances to keep

After a certain number of job runs for a given job, older NetApp ONTAP objects for that job are purged from the Inventory. Enter the number of job runs for which high-level NetApp ONTAP objects are to be retained.

7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

8. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

9. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

Note: If you selected the **Start job now** option, the job runs.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Create an Inventory Job Definition - NetApp ONTAP File](#) on page **186**
- [Edit a Job Definition](#) on page **342**

- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**

Create an Inventory Job Definition - NetApp ONTAP File

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

A NetApp ONTAP File Inventory job provides the framework to catalog and collect information about low-level NetApp ONTAP objects. You can select one or more NetApp ONTAP providers in a single job definition for cataloging. NetApp ONTAP 7-Mode and Cluster-Mode are both supported.

During cataloging, properties for NetApp ONTAP files are collected and stored. Run a NetApp ONTAP File Inventory job to analyze your low-level NetApp ONTAP environment in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer snapshot sprawl, storage overutilization, and other storage inefficiencies.

Note: For cataloging high-level NetApp ONTAP objects, create a NetApp ONTAP Storage Inventory job definition.

Tip: NetApp ONTAP File Inventory jobs may take a long time to run because they catalog at a file level. Consider creating NetApp ONTAP File Inventory job definitions that constrain the number of file system objects to catalog by limiting the number of storage systems and volume processed by a single job. You can also stagger the scheduled run times of the jobs to prevent them from running concurrently. Consider running NetApp ONTAP File Inventory jobs less frequently than you run NetApp ONTAP Storage Inventory jobs.

BEFORE YOU BEGIN:

- At least one NetApp ONTAP provider must be associated with a NetApp ONTAP File Inventory job definition. Before defining an Inventory job, add NetApp ONTAP providers. See [Register a Provider](#) on page 75.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

Best Practice: Create a schedule before creating a job definition so that you can easily add the schedule to the job definition.

To create a NetApp ONTAP File Inventory job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **NetApp ONTAP**.
2. Click **New** , then select **Inventory NetApp ONTAP Files**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.

4. From the list of available sites, select one or more volumes to catalog. For Cluster-Mode providers, the SVM name appears in parentheses after the volume name. To view the number of files on the selected volume, hover your cursor over the volume name.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

Skip root volume

Select this option to avoid cataloging the root volume of your NetApp ONTAP objects.

Skip unsupported volumes

Select this option to avoid cataloging unsupported volumes such as volumes that are offline. Selecting this option also avoids cataloging volumes with i2p disabled when **Catalog all available snapshots** is set to Yes and **Traversal Method** is set to Snapdiff.

Connection timeout (secs)

To run a catalog job, the application needs to connect with the resource. If there is no response within a certain time limit, it times out and the job session fails. Enter the number of seconds to wait before timing out.

Number of catalog instances to keep

After a certain number of job runs for a given job, older low-level NetApp ONTAP objects for that job are purged from the Inventory. Enter the number of job runs for which low-level NetApp ONTAP objects are to be retained. Note that by default, the old data for the job is purged from the IBM Spectrum Copy Data Management Inventory after the newer data is cataloged.

Condense catalog before run

Select this option to purge older low-level NetApp ONTAP objects from the IBM Spectrum Copy Data Management Inventory before newer data is cataloged through a NetApp ONTAP File Inventory job.

Condense catalog after failed run

Select this option to purge NetApp ONTAP objects from the IBM Spectrum Copy Data Management Inventory for the unsuccessful run of the NetApp ONTAP File Inventory job.

Traversal Method

This option indicates the methodology to employ when cataloging snapshots. IBM Spectrum Copy Data Management honors your preference if it is supported for the particular system configuration. If the selected preference is not supported for your system configuration, the operation fails.

SnapDiff. IBM Spectrum Copy Data Management performs cataloging based on snapshot differences. If SnapDiff is selected, the default number of files requested in each query is 256 and the default maximum number of volumes that are simultaneously cataloged is 8.

Filewalk. IBM Spectrum Copy Data Management retrieves owner information for files and folders for volumes that have CIFS shares. This option is used in conjunction with the IBM Spectrum Copy Data Management Filewalker tool. Once cataloging completes, searchable owner information is added to the Inventory. See the Filewalker documentation for more information.

Base snapshot for catalog

For file cataloging, a base snapshot on the storage system is required.

Create a new snapshot before cataloging. A new snapshot is created when the job session begins and is deleted after cataloging. This is the preferred method if sufficient resources are available.

Use latest snapshot. A new snapshot is not created; instead the latest healthy snapshot available for that volume is used to assist with file level cataloging. You can exclude certain snapshots from the set of snapshots IBM Spectrum Copy Data Management selects from by entering character patterns for those snapshots. Use a comma to separate each pattern.

Note: To catalog a volume that does not support snapshot creation, such as a SnapMirror volume, select **Use latest snapshot**.

Catalog all available snapshots

Yes. Catalogs snapshots on the volume in addition to the base snapshot. IBM Spectrum Copy Data Management examines additional snapshots ensuring the Inventory has the latest data. You can exclude certain snapshots from the set of snapshots IBM Spectrum Copy Data Management catalogs by entering character patterns for those snapshots. Use a comma to separate each pattern.

No. Catalogs only the base snapshot.

Note: By cataloging all available snapshots, you can view multiple versions of your files through a file-level search. Additional versions are available on the file's properties pane.

7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

8. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

9. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

Note: If you selected the **Start job now** option, the job runs.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Create an Inventory Job Definition - NetApp ONTAP Storage](#) on page **183**
- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**
- [Search and Filter Guidelines](#) on page **548**

Create an Inventory Job Definition - HPE Nimble Storage

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

A HPE Nimble Storage Inventory job provides the framework to catalog and collect information about high-level objects on your HPE Nimble Storage systems. You can select one or more HPE Nimble Storage providers in a single job definition for cataloging.

Run a HPE Nimble Storage Inventory job to analyze your HPE Nimble Storage environment in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer sprawl, storage overutilization, and other storage inefficiencies.

BEFORE YOU BEGIN:

- At least one HPE Nimble Storage provider must be associated with a HPE Nimble Storage Inventory job definition. Before defining an Inventory job, add HPE Nimble

Storage providers. See [Register a Provider](#) on page 75.

- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create a HPE Nimble Storage Inventory job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **HPE Nimble Storage**.
2. Click **New** , then select **Inventory**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more providers to catalog.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

Number of catalog instances to keep

After a certain number of job runs for a given job, older HPE Nimble Storage objects for that job are purged from the Inventory. Enter the number of job runs for which high-level HPE Nimble Storage objects are to be retained.

7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

8. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

9. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

Note: If you selected the **Start job now** option, the job runs.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**

Create an Inventory Job Definition - Pure Storage FlashArray

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

A Pure Storage FlashArray Inventory job provides the framework to catalog and collect information about high-level objects on your Pure Storage systems. You can select one or more Pure Storage providers in a single job definition for cataloging.

Run a Pure Storage FlashArray Inventory job to analyze your Pure Storage environment in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer sprawl, storage overutilization, and other storage inefficiencies.

BEFORE YOU BEGIN:

- At least one Pure Storage provider must be associated with a Pure Storage FlashArray Inventory job definition. Before defining an Inventory job, add Pure Storage providers. See [Register a Provider](#) on page 75.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create a Pure Storage FlashArray Inventory job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **Pure Storage FlashArray**.
2. Click **New** , then select **Inventory**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more providers to catalog.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

Number of catalog instances to keep

After a certain number of job runs for a given job, older Pure Storage objects for that job are purged from the Inventory. Enter the number of job runs for which high-level Pure Storage objects are to be retained.

7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

8. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

9. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

Note: If you selected the **Start job now** option, the job runs.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**

Create an Inventory Job Definition - VMware

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

A VMware Inventory job provides the framework to catalog and collect information about VMware objects. You can select one or more VMware providers in a single job definition for cataloging.

Run a VMware Inventory job to analyze your VMware environment in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer snapshot sprawl, storage overutilization, and other storage inefficiencies.

BEFORE YOU BEGIN:

- At least one VMware provider must be associated with a VMware Inventory job definition. Before defining an Inventory job, add VMware providers. See [Register a Provider](#) on page 75.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create a VMware Inventory job definition:

1. Click the **Jobs**  tab. Expand the **Hypervisor** folder, then select **VMware**.
2. Click **New** , then select **Inventory**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more providers to catalog.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

Connection timeout (secs)

To run a catalog job, the application needs to connect with the resource. If there is no response within a certain time limit, it times out and the job session fails. Enter the number of seconds to wait before timing out.

Number of catalog instances to keep

After a certain number of job runs for a given job, older VMware objects for that job are purged from the Inventory. Enter the number of job runs for which VMware objects are to be retained.

7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

8. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

9. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

Note: If you selected the **Start job now** option, the job runs.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page 162.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page 164.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Edit a Job Definition](#) on page 342
- [Delete a Job Definition](#) on page 343
- [Create a Schedule](#) on page 155

Backup Jobs

The topics in the following section cover Backup job definitions as well as creating VMware Backup job proxies.

Create a Backup Job Definition - InterSystems Caché

IBM Spectrum Copy Data Management provides application database copy management through application-consistent backup creation, cloning, and recovery. IBM Spectrum Copy Data Management copy management leverages the snapshot and replication features of the underlying storage platform to create, replicate, clone, and restore backups of InterSystems Caché instances.

IBM Spectrum Copy Data Management auto-discovers databases and enables backups only of eligible instances. To be eligible for backup, instances must reside on supported storage platforms.

BEFORE YOU BEGIN:

- Create and run a Database Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - Database](#) on page 170.
- Review InterSystems Caché requirements. See [InterSystems Caché Requirements](#) on page 34.
- Configure an SLA Policy. See [Configure SLA Policies](#) on page 133.
- Review InterSystems Caché requirements. See [InterSystems Caché Requirements](#) on page 34.

CONSIDERATIONS:

- InterSystems Caché backup jobs occur at the instance level.
- Note that it is possible to scan in an InterSystems Caché backup failover member instance or an async member instance and run snapshots against the mirror copy instead of the primary failover member.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create an InterSystems Caché Backup job definition:

1. Click the **Jobs**  tab. Expand the **Database** folder, then select **InterSystems Caché**.
2. Click **New** , then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more resources to back up.
5. Select an SLA Policy that meets your backup data criteria.

- Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs**  tab. Repeat as necessary to add additional SLA Policies to the job definition.
If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.
Note: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.
- To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.
- To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Skip IA Mount points and/or databases

Enable to skip Instant Disk Restore objects. By default, this option is enabled.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

- Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

10. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Create an InterSystems Caché Restore job definition. See [Create a Restore Job Definition - InterSystems Caché](#) on page **252**.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**
- [Create a Restore Job Definition - InterSystems Caché](#) on page **252**

Create a Backup Job Definition - SAP HANA

IBM Spectrum Copy Data Management provides application database copy management through application-consistent backup creation, cloning, and recovery. IBM Spectrum Copy Data Management copy management leverages the snapshot and replication features of the underlying storage platform to create, replicate, clone, and restore backups of SAP HANA databases.

IBM Spectrum Copy Data Management auto-discovers databases and enables backups only of eligible databases. To be eligible for backup, application databases must reside on supported storage platforms.

Log Backup - The log backup feature enables continuous backup of Archive logs to a specified destination. IBM Spectrum Copy Data Management leverages archived logs to enable point-in-time recoveries of databases to facilitate RPOs.

BEFORE YOU BEGIN:

- Review SAP HANA requirements. See [SAP HANA Requirements](#) on page **51**.
- Review SAP HANA prerequisites. See [Prerequisites](#) on page **52**.
- Create and run a Database Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - Database](#) on page **170**.
- Configure an SLA Policy. See [Configure SLA Policies](#) on page **133**.

CONSIDERATIONS:

- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page **75**.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page **155**.

To create a SAP HANA Backup job definition:

1. Click the **Jobs**  tab. Expand the **Database** folder, then select **SAP HANA**.
2. Click **New** , then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more providers to back up.

Note: You cannot select a database if it is not eligible for protection. Hover your cursor over the database name to view the reasons the database is ineligible, such as the database files, control files, or redo log files are stored on unsupported storage.

5. Select an SLA Policy that meets your backup data criteria.

- Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs**  tab. Repeat as necessary to add additional SLA Policies to the job definition.
If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.
Note: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.
- To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.
- To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Skip IA Mount points and/or databases

Enable to skip Instant Disk Restore objects. By default, this option is enabled.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

- Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

10. To edit Log Backup options before creating the job definition, click **Log Backup**. If **Backup Logs** is selected, IBM Spectrum Copy Data Management backs up database logs then protects the underlying disks. Select resources in the **Select resource(s) to add archive log destination** field. Database logs are backed up to the directory entered in the **Use Universal Destination Mount Point** field, or in the Mount Point field after resources are selected. The destination must already exist, must reside on storage from a supported vendor, and the SAP HANA user needs to have full read and write access. For more information, see the [Prerequisites](#) on page 52.

If multiple databases are selected for backup, then each of the servers hosting the databases must have their Destination Mount Points set individually. For example, if two databases, one from Server A and one from Server B, are added to the same job definition, and a single mount point named `/logbackup` is defined in the job definition, then you must create separate disks for each server and mount them both to `/logbackup` on the individual servers. When the mount point is changed, you must manually go in and clean up the previous log backup directory path.

To disable a log backup schedule on the SAP HANA server, edit the associated SAP HANA Backup job definition and deselect the checkbox next to the database on which you wish to disable the log backup schedule in the **Select resource(s) for log backup destination** field, then save and re-run the job. When the mount point is disabled, you must manually go in and clean up the log backup directory path.

Note: The job definition must be saved and re-run for mount point changes or disablement to take effect.

The default setting for pruning SAP HANA log backups is 7 days. This value may be adjusted in the property file located in `/opt/virgo/repository/ecx-`

`usr/com.syncsort.dp.xsb.serviceprovider.properties`. Modify the `application.logpurge.days` parameter to the desired value. Finally, restart the virgo service by issuing the following command:

```
systemctl restart virgo.service
```

11. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page 162.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page 164.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

- Create a SAP HANA Restore job definition. See [Create a Restore Job Definition - SAP HANA](#) on page **257**.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**
- [Create a Restore Job Definition - SAP HANA](#) on page **257**

Create a Backup Job Definition - Oracle

IBM Spectrum Copy Data Management provides application database copy management through application-consistent backup creation, cloning, and recovery. IBM Spectrum Copy Data Management copy management leverages the snapshot and replication features of the underlying storage platform to create, replicate, clone, and restore backups of Oracle databases. Archive log destinations as well as universal destination mount points are supported. Archived logs are automatically deleted upon reaching defined retention.

IBM Spectrum Copy Data Management auto-discovers databases and enables backups only of eligible databases. To be eligible for backup, application databases must reside on supported storage platforms.

The following options are available for Oracle Backup jobs:

RMAN Integration - Oracle Recovery Manager (RMAN), a command-line and Enterprise Manager-based tool, is the method preferred by Oracle database administrators for backup and recovery of Oracle databases, including maintaining an RMAN repository. The retention of RMAN cataloged data is managed by settings in Oracle. IBM Spectrum Copy Data Management automates cataloging of Oracle database backups in the RMAN recovery catalog, enabling database administrators to leverage RMAN for verification and advanced recovery.

Data Masking - Data masking is used to hide confidential data by replacing it with fictitious data. This feature is used when making data copies for DevTest or other use cases.

Log Backup - The log backup feature enables continuous backups of Archive logs to a specified destination. Archive log retention is managed by settings in RMAN. IBM Spectrum Copy Data Management leverages archived logs to enable point-in-time recoveries of databases to facilitate RPOs.

BEFORE YOU BEGIN:

- Create and run a Database Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - Database](#) on page 170.
- Configure an SLA Policy. See [Configure SLA Policies](#) on page 133.
- Review Oracle requirements. See [Oracle Requirements](#) on page 41 and [Oracle Database Support FAQ](#) on page 563.

ORACLE DATABASE CONSIDERATIONS:

- To ensure that filesystem permissions are retained correctly when IBM Spectrum Copy Data Management moves Oracle data between servers, ensure that the user and group IDs of the Oracle users (e.g. oracle, oinstall, dba) are consistent across all the servers. Refer to Oracle documentation for recommended uid and gid values.
- If Oracle data resides on LVM volumes, you must stop and disable the lvm2-lvmetad service before running Backup or Restore jobs. Leaving the service enabled can prevent volume groups from being resigatured correctly during restore and can lead to data corruption if the original volume group is also present on the same system. To disable the lvm2-lvmetad service, run the following commands:

```
systemctl stop lvm2-lvmetad
systemctl disable lvm2-lvmetad
```

Next, disable `lvmetad` in the LVM config file. Edit the file `/etc/lvm/lvm.conf` and set:

```
use_lvmetad = 0
```

- Note that Oracle databases must be registered in the recovery catalog before running an Oracle Backup job utilizing the **Record copies in RMAN recovery catalog** feature.
- In your Linux environment, if Oracle data or logs reside on LVM volumes, ensure the LVM version is 2.0.2.118 or later.
- For Oracle 12c databases, backups are created without placing the database in hot backup mode through Oracle Storage Snapshot Optimization. All associated snapshot functionality is supported. This feature requires the *Advanced Compression* feature of Oracle to be licensed. If this feature is not licensed in your environment, perform the following procedure to disable Snapshot Optimization and force the use of hot backup mode:

Create the file `/etc/guestapps.conf` on the Oracle server and add the following to it:

```
[DEFAULT]
```

```
skipSnapshotOptimization = true
```

If the file already exists, edit it and add the parameter under the existing `[DEFAULT]` section. This is a per-host setting. The parameter must be set in this file on each Oracle server where you want to force the use of hot backup mode.

- NOARCHIVELOG databases are not eligible for point-in-time recovery. NOARCHIVELOG databases can only be recovered to specific or latest versions. If upgrading from previous versions of IBM Spectrum Copy Data Management, the associated Oracle Inventory job must be re-run after upgrading to discover NOARCHIVELOG databases.
- When the option to create an additional log destination is selected, IBM Spectrum Copy Data Management automatically purges the logs under this new location after each successful backup. For IBM SVC, IBM Spectrum Copy Data Management purges logs after a FlashCopy operation but not after a Global Mirror operation. If both FlashCopy and Global Mirror are enabled for a database (whether in separate job definitions or the same), IBM Spectrum Copy Data Management purges the logs after the FlashCopy operation only. For databases that are protected only by a Global Mirror workflow, IBM Spectrum Copy Data Management does not purge the logs at all so they must be deleted using a retention policy externally managed by a database administrator, for example,

using RMAN. Note that in any case, IBM Spectrum Copy Data Management does not purge logs from other log destinations so they must also be externally managed.

- If an Oracle Inventory job runs at the same time or short period after an Oracle Backup job runs, copy errors may occur due to temporary mounts that are created during the Backup job. As a best practice, schedule Oracle Inventory jobs so that they do not overlap with Oracle Backup jobs.

CONSIDERATIONS:

- Note that point-in-time recovery is not supported when one or more datafiles are added to the database in the period between the chosen point-in-time and the time that the preceding Backup job ran.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create an Oracle Backup job definition:

1. Click the **Jobs**  tab. Expand the **Database** folder, then select **Oracle**.
2. Click **New** , then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more providers to back up. Expand Oracle home directories to view associated application databases.

Note: You cannot select a database if it is not eligible for protection. Hover your cursor over the database name to view the reasons the database is ineligible, such as the database files, control files, or redo log files are stored on unsupported storage.

5. Select an SLA Policy that meets your backup data criteria.
6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs**  tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Note: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum Concurrent Tasks

Set the maximum amount of concurrent transfers between the source and the destination.

Skip IA Mount points and/or databases

Enable to skip Instant Disk Restore objects. By default, this option is enabled.

Record copies in RMAN local repository

Enable to create a local backup of the Recovery Manager (RMAN) catalog during the running of Oracle Backup job. RMAN catalogs can be used for backup, recovery, and maintenance of Oracle databases outside of IBM Spectrum Copy Data Management.

Record copies in RMAN recovery catalog

If Record copies in RMAN local repository is selected, select Record copies in RMAN recovery catalog to also create a remote RMAN catalog. Select an eligible Remote Catalog Database from the list of available sites. Select a Recovery Catalog Owner from the list of available Identities, or create a new Recovery Catalog Owner, then click **OK**.

Note that Oracle databases must be registered in the recovery catalog before running an Oracle Backup job utilizing the **Record copies in RMAN recovery catalog** feature.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands,

such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

10. To edit Log Backup options before creating the job definition, click **Log Backup**. If **Create additional archive log destination** is selected, IBM Spectrum Copy Data Management backs up database logs then protects the underlying disks. Select resources in the **Select resource(s) to add archive log destination** field. Database logs are backed up to the directory entered in the **Universal destination directory** field, or in the **Directory** field after resources are selected. The destination must already exist and must reside on storage from a supported vendor.

The default option is **Use existing archive log destination(s)**. Note that IBM Spectrum Copy Data Management automatically discovers the location where Oracle writes archived logs. If this location resides on storage from a supported vendor, IBM Spectrum Copy Data Management can protect it. If the existing location is not on supported storage, or if you wish to create an additional backup of database logs, enable the **Create additional archive log destination** option, then specify a path that resides on supported storage. When enabled, IBM Spectrum Copy Data Management configures the database to start writing

archived logs to this new location in addition to any existing locations where the database is already writing logs.

Note: NOARCHIVELOG databases are not eligible for log backup as they do not have archive logging enabled.

If multiple databases are selected for backup, then each of the servers hosting the databases must have their destination directories set individually. For example, if two databases from Server A and Server B are added to the same job definition, and a single destination directory named /logbackup is defined in the job definition, then you must create separate disks for both servers and mount them both to /logbackup on the individual servers.

If the **No archive logs / Use existing archive log destination(s)** option is selected, IBM Spectrum Copy Data Management does not automatically purge any archived logs. The retention of archived logs must be managed externally, for example using RMAN. In order to support point-in-time recovery, ensure that the retention period is at least large enough to retain all archived logs between successive runs of the Oracle Backup job.

If the **Create additional archive log destination** option is selected, IBM Spectrum Copy Data Management automatically manages the retention of only those archived logs that are under the new destination specified in the job definition. After a successful backup, logs older than that backup are automatically deleted from the IBM Spectrum Copy Data Management-managed destination. Even in this case, IBM Spectrum Copy Data Management does not control the deletion of archived logs in other pre-existing destinations so they must still be managed externally as described above.

If the **Create additional archive log destination** option is selected, IBM Spectrum Copy Data Management makes a one-time configuration change to the database to add the specified location as a parameter `log_archive_dest_<num>` in the database's archive log destinations. If you delete the IBM Spectrum Copy Data Management job definition, the database parameter is not affected so if you want to stop using the log destination, you may need to manually disable it this parameter.

11. To edit Data Masking options before creating the job definition, click **Data Masking**. If enabled, IBM Spectrum Copy Data Management mounts snapshot copies of the protected database onto a user-specified staging server or source server. Select resources to be masked from the list of available databases, select a backup to mask, and an Oracle home where masking takes place. Set a trigger, then in the **Enter path to masking command on Oracle Server** field, enter the full path to an external script or tool to perform the data masking. For example, /home/oracle/tools/maskDatabase.sh.

Whether the masking takes place on the source server or a staging server, the masking process spins up the temporary database with a unique random name (for example, mask1234) that does not conflict with any other instance on that system. IBM Spectrum Copy Data Management then invokes the masking script with three arguments: the Oracle Home path, the new instance name, and the original instance name. For example:

```
/path/to/masking/script /u01/app/home1 mask1234 proddb
```

IBM Spectrum Copy Data Management spins up a clone of the database, then executes the user-specified command to perform masking. When the command completes successfully, IBM Spectrum Copy Data Management cleans up the clone database, and catalogs and saves the masked copies which are then available for selection in the DevOps workflow of IBM Spectrum Copy Data Management Restore jobs.

Note: User-defined masking scripts that were in existence before IBM Spectrum Copy Data Management 2.2.7.4 must be updated to ensure they read the correct arguments and connect to the appropriate instance to perform masking.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Create an Oracle Restore job definition. See [Create a Restore Job Definition - Oracle](#) on page **264**.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**
- [Create a Restore Job Definition - Oracle](#) on page **264**

Create a Backup Job Definition - SQL

IBM Spectrum Copy Data Management provides application database copy management through application-consistent backup creation, cloning, and recovery. IBM Spectrum Copy Data Management copy management leverages the snapshot and replication features of the underlying storage platform to create, replicate, clone, and restore copies of Microsoft SQL Servers. Archive log destinations as well as universal destination mount points are supported. Archived logs are automatically deleted upon reaching defined retention.

IBM Spectrum Copy Data Management auto-discovers databases and enables backups only of eligible databases. To be eligible for backup, application databases must reside on supported storage platforms.

The following options are available for SQL Backup jobs:

Log Backup - The log backup feature enables continuous backup of Archive logs to a specified destination. IBM Spectrum Copy Data Management leverages archived logs to enable point-in-time recoveries of databases to facilitate RPOs.

BEFORE YOU BEGIN:

- Create and run a Database Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - Database](#) on page 170.
- Configure an SLA Policy. See [Configure SLA Policies](#) on page 133.
- Review SQL requirements. See [Microsoft SQL Server Requirements](#) on page 57 and [Microsoft SQL Server Support FAQ](#) on page 584.

MICROSOFT SQL SERVER CONSIDERATIONS:

- UUID must be enabled to perform Microsoft SQL-based backup functions. To enable, power off the guest machine through the vSphere client, then select the guest and click **Edit Settings**. Select **Options**, then **General** under the Advanced section. Select **Configuration Parameters...**, then find the `disk.EnableUUID` parameter. If set to FALSE, change the value to TRUE. If the parameter is not available, add it by clicking **Add Row**, set the value to TRUE, then power on the guest.
- You must manually create a directory to store VSS provider logs when running IBM Spectrum Copy Data Management 2.2.6 and earlier. Create the following directory structure on the SQL server: `c:\temp\CDM\logs`
- During a Backup job, an error may display in the logs stating that the size of the manifest file is a specific size, and the snapshot is considered not application consistent. If this error displays, it may be due to the number of disks associated with a particular SCSI controller. As a best practice, configure your disks so that all controllers have at least half of their slots empty.

- IBM Spectrum Copy Data Management supports one Microsoft SQL database per SQL Backup job, therefore you must avoid protecting a SQL database through multiple Backup jobs.
- Note that IBM Spectrum Copy Data Management does not support log backup of Simple recovery models.
- An AlwaysOn of a replica of a SQL cluster instance is not supported. Replicas are limited standalone SQL servers and instances.

CONSIDERATIONS:

- Note that point-in-time recovery is not supported when one or more datafiles are added to the database in the period between the chosen point-in-time and the time that the preceding Backup job ran.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create a SQL Backup job definition:

1. Click the **Jobs**  tab. Expand the **Database** folder, then select **SQL**.
2. Click **New** , then select **Backup**. The job editor opens.
3. Select a **Standalone or Failover Cluster** or **Always On Availability Group** workflow template.
4. Enter a name for your job definition and a meaningful description.
5. From the list of available sites select a provider to back up. Expand servers to view associated application databases.

Note: You cannot select a database if it is not eligible for protection. Hover your cursor over the database name to view the reasons the database is ineligible, such as the database files, control files, or redo log files are stored on unsupported storage.

6. Select an SLA Policy that meets your backup data criteria.
7. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs**  tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Note: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

- To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.
- To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum Concurrent Tasks

Set the maximum amount of concurrent transfers between the source and the destination.

Skip IA Mount points and/or databases

Enable to skip Instant Disk Restore objects. By default, this option is enabled.

Maximum Concurrent Snapshots on ESX

Set the maximum number of concurrent snapshots on the vCenter.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

11. To edit Log Backup options before creating the job definition, click **Log Backup**. If **Backup Logs** is selected, IBM Spectrum Copy Data Management backs up database logs then protects the underlying disks. Select resources in the **Select resource(s) to add archive log destination** field. Database logs are backed up to the directory entered in the **Use Universal Destination Mount Point** field, or in the Mount Point field after resources are selected. The destination must already exist and must reside on storage from a supported vendor.

Note: Always On job definitions must specify the log destination as a path in the following format: \\server\share\optional_subfolder. The server can be either an IP address or hostname that is resolvable from the IBM Spectrum Copy Data Management appliance.

If multiple databases are selected for backup, then each of the servers hosting the databases must have their Destination Mount Points set individually. For example, if two databases, one from Server A and one from Server B, are added to the same job definition, and a single mount point named /logbackup is defined in the job definition, then you must create separate disks for each server and mount them both to /logbackup on the individual servers.

IBM Spectrum Copy Data Management automatically truncates post log backups of databases that it backs up. If database logs are not backed up with IBM Spectrum Copy Data Management, logs are not truncated by IBM Spectrum Copy Data Management and must be managed separately.

To disable a log backup schedule on the SQL server, edit the associated SQL Backup job definition and deselect the checkbox next to the database on which you wish to disable the log backup schedule in the **Select resource(s) for log backup destination** field, then save and re-run the job. Note that the job definition must be saved and re-run for the disablement to take effect.

When SQL backup job completes with log backups enabled, all transaction logs up to the point of the job completing are purged from the SQL server. Note that log purging will only occur if the SQL Backup job completes successfully. If log backups are disabled during a re-run of the job, log purging will not occur.

If a source database is overwritten, all old transaction logs up to that point are placed in a “condense” directory once the restoration of the original database completes. When the next run of the SQL Backup job completes, the contents of the condense folder is removed.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Create a SQL Application Restore job definition. See [Create a Restore Job Definition - SQL](#) on page **270**.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**
- [Create a Restore Job Definition - SQL](#) on page **270**

Create a Backup Job Definition - File System

IBM Spectrum Copy Data Management provides application database copy management of physical Windows, Linux, and AIX file systems through a File System Backup job.

BEFORE YOU BEGIN:

- Review File System requirements. See [File System Requirements](#) on page 28.
- Create and run a File System Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - File System](#) on page 173.
- Configure an SLA Policy. See [Configure SLA Policies](#) on page 133.

CONSIDERATIONS:

- An AlwaysOn of a replica of a SQL cluster instance is not supported. Replicas are limited standalone SQL servers and instances.
- Note that point-in-time recovery is not supported when one or more datafiles are added to the database in the period between the chosen point-in-time and the time that the preceding Backup job ran.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create a File System Backup job:

1. Click the **Jobs**  tab. Expand the **Database** folder, then select **File System**.
2. Click **New** , then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more providers to back up.

Note: You cannot select a database if it is not eligible for protection. Hover your cursor over the database name to view the reasons the database is ineligible, such as the database files, control files, or redo log files are stored on unsupported storage.

5. Select an SLA Policy that meets your backup data criteria.
6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs**  tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Note: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum Concurrent Tasks

Set the maximum amount of concurrent transfers between the source and the destination.

Skip IA Mount points and/or databases

Enable to skip Instant Disk Restore objects. By default, this option is enabled.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

Note: If adding a script to a Windows-based File System job definition, the user running the script must have the "Log on as a service" right enabled, which is required for running prescripts and postscripts. For more information about the "Log on as a service" right, see <https://technet.microsoft.com/en-us/library/cc794944.aspx>.

- Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

- When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page 162.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page 164.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Create an OS Volume Restore job definition. See [Create a Restore Job Definition - File System](#) on page 277.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**
- [Create a Restore Job Definition - File System](#) on page **277**

Create a Backup Job Definition - DellEMC Unity

Back up DellEMC Unity data with snapshots, replication, and VM Copies using a DellEMC Unity Backup job. The RPO and copy data parameters are defined in an SLA Policy, which is then applied to the Backup job definition along with a specified activation time to meet your copy data criteria.

BEFORE YOU BEGIN:

- Create and run a DellEMC Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - DellEMC Unity](#) on page 175.
- Configure an SLA Policy. See [Configure SLA Policies](#) on page 133.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- Note that before running replication jobs, replication connections must be established between VNX arrays. Create replication connections through the DellEMC Unisphere wizard found under **Hosts > Replication Connections**.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create a DellEMC Backup job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **DellEMC Unity**.
2. Click **New** , then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the drop-down menu select **LUNs** or **File Systems**. From the list of available sites, select one or more providers to back up.
5. Select an SLA Policy that meets your backup data criteria.
6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs**  tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Note: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.

- To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum Concurrent Tasks

Set the maximum amount of concurrent transfers between the source and the destination.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

- Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

- When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page 162.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page 164.

- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Use the Inventory Browse feature to review the recovery point. See [Browse Inventory](#) on page **357**.
- Create a DellEMC Unity Restore job definition. See [Create a Restore Job Definition - DellEMC Unity](#) on page **281**.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**
- [Create a Restore Job Definition - DellEMC Unity](#) on page **281**

Create a Backup Job Definition - IBM Spectrum Accelerate

Back up IBM Spectrum Accelerate data with FlashCopies, Global Mirrors, and VM Copies using an IBM Spectrum Accelerate Backup job. The RPO and copy data parameters are defined in an SLA Policy, which is then applied to the Backup job definition, along with a specified activation time to meet your copy data criteria.

BEFORE YOU BEGIN:

- Create and run an IBM Spectrum Accelerate Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - IBM Spectrum Accelerate](#) on page 177.
- Configure an SLA Policy. See [Configure SLA Policies](#) on page 133.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- IBM providers utilize port 22 for communication with IBM Spectrum Copy Data Management.
- Note that snapshot postscript functionality applies only to FlashCopy subpolicies.
- In IBM storage environments, port grouping and IP partnerships are required to enable remote copy connections. See IBM's [SAN Volume Controller and Storwize Family Native IP Replication Guide](#).
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create an IBM Spectrum Accelerate Backup job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **IBM Spectrum Accelerate**.
2. Click **New** , then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more resources to back up
5. Select an SLA Policy that meets your backup data criteria.
6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs**  tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Note: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum Concurrent Tasks

Set the maximum amount of concurrent transfers between the source and the destination.

Create Consistency Group

If multiple volumes are selected in the Source tab (for example, volumes that contain data tied to an application) enable this option to add the volumes to a FlashCopy or Global Mirror Consistency Group to perform Copy Data functions on the entire group. If the associated SLA Policy contains both FlashCopy and Global Mirror subpolicies, a separate Consistency Group will be created for each copy type. Note that if more than one IBM provider is selected in the job definition, a Consistency Group will be created for each provider. Consistency Groups are named based on the prefix provided during job creation plus the job name.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands,

such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page **152**.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page **334**.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page **553**.

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

10. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

- Use the Inventory Browse feature to review the recovery point. See [Browse Inventory](#) on page **357**.
- Create an IBM Spectrum Accelerate Restore job definition. See [Create a Restore Job Definition - IBM Spectrum Accelerate](#) on page **287**.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**
- [Create a Restore Job Definition - IBM Spectrum Accelerate](#) on page **287**

Create a Backup Job Definition - IBM Spectrum Virtualize

Back up IBM data with FlashCopies, Global Mirrors, and VM Copies using an IBM Spectrum Virtualize Backup job. The RPO and copy data parameters are defined in an SLA Policy, which is then applied to the Backup job definition, along with a specified activation time to meet your copy data criteria. Supported sources include IBM Spectrum Virtualize storage systems.

BEFORE YOU BEGIN:

- Create and run an IBM Spectrum Virtualize Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - IBM Spectrum Virtualize](#) on page **181**.
- Configure an SLA Policy. See [Configure SLA Policies](#) on page **133**.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page **75**.

CONSIDERATIONS:

- IBM providers utilize port 22 for communication with IBM Spectrum Copy Data Management.
- Note that snapshot postscript functionality applies only to FlashCopy subpolicies.
- In IBM storage environments, port grouping and IP partnerships are required to enable remote copy connections. See IBM's [SAN Volume Controller and Storwize Family Native IP Replication Guide](#).
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page **155**.

To create an IBM Spectrum Virtualize Backup job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **IBM Spectrum Virtualize**.
2. Click **New** , then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more providers to back up. To exclude Flash Copies from the list of sources, select **Exclude Flash Copies**.
5. Select an SLA Policy that meets your backup data criteria.
6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs**  tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Note: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Skip the Flash Copy Target Volumes

Select this option to ensure FlashCopy target volumes are excluded from jobs associated with the SLA Policy.

Maximum Concurrent Tasks

Set the maximum amount of concurrent transfers between the source and the destination.

Create Consistency Group

If multiple volumes are selected in the Source tab (for example, volumes that contain data tied to an application) enable this option to add the volumes to a FlashCopy or Global Mirror Consistency Group to perform Copy Data functions on the entire group. If the associated SLA Policy contains both FlashCopy and Global Mirror subpolicies, a separate Consistency Group will be created for each copy type. Note that if more than one IBM provider is selected in the job definition, a Consistency Group will be created for each provider. Consistency Groups are named based on the prefix provided during job creation plus the job name.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

10. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Use the Inventory Browse feature to review the recovery point. See [Browse Inventory](#) on page **357**.
- Create an IBM Spectrum Virtualize Restore job definition. See [Create a Restore Job Definition - IBM Spectrum Virtualize](#) on page **293**.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**
- [Create a Restore Job Definition - IBM Spectrum Virtualize](#) on page **293**

Create a Backup Job Definition - NetApp ONTAP

Back up NetApp ONTAP data using a NetApp ONTAP Backup job. A NetApp ONTAP Backup job consists of snapshot, mirror, and vault sub-policies defined in an SLA Policy, each with their own set of options to give you more control of your NetApp ONTAP protection needs. After an initial primary snapshot is added to the workflow, additional vaults and mirrors ensure your data is replicated to multiple locations.

BEFORE YOU BEGIN:

- Create and run a NetApp ONTAP Storage Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - NetApp ONTAP Storage](#) on page **183**.
- Configure an SLA Policy. See [Configure SLA Policies](#) on page **133**.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page **75**.

CONSIDERATIONS:

- Note that NetApp ONTAP Backup jobs can only vault or mirror snapshots created through IBM Spectrum Copy Data Management jobs.
- Note that cloned volumes will not be replicated through a Backup job.
- Note that snapshot postscript functionality applies only to NetApp ONTAP storage snapshot subpolicies.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page **155**.

To create a NetApp ONTAP Backup job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **NetApp ONTAP**.
2. Click **New** , then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more resources to back up.
5. Select an SLA Policy that meets your backup data criteria.
6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs**  tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Note: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum Concurrent Tasks

Set the maximum amount of concurrent transfers between the source and the destination.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by

entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page **334**.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page **553**.

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

- Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

- When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Use the Inventory Browse feature to review the recovery point. See [Browse Inventory](#) on page **357**.
- Create a NetApp ONTAP Restore job definition. See [Create a Restore Job Definition - NetApp ONTAP](#) on page **299**.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**

- [Create a Schedule](#) on page **155**
- [Create a Restore Job Definition - NetApp ONTAP](#) on page **299**

Create a Backup Job Definition - HPE Nimble Storage

Back up HPE Nimble Storage data with snapshots using a HPE Nimble Storage Backup job. The RPO and copy data parameters are defined in an SLA Policy, which is then applied to the Backup job definition along with a specified activation time to meet your copy data criteria.

BEFORE YOU BEGIN:

- Create and run a HPE Nimble Storage Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - HPE Nimble Storage](#) on page **189**.
- Configure an SLA Policy. See [Configure SLA Policies](#) on page **133**.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page **75**.

CONSIDERATIONS:

- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page **155**.

To create a HPE Nimble Storage Backup job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **HPE Nimble Storage**.
2. Click **New** , then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more resources to back up.
5. Select an SLA Policy that meets your backup data criteria.
6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs**  tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Note: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum Concurrent Tasks

Set the maximum amount of concurrent transfers between the source and the destination.

Create Consistency Group

If multiple volumes are selected in the Source tab (for example, volumes that contain data tied to an application) enable this option to add the volumes to a Consistency Group to perform backup functions on the entire group. If the associated SLA Policy contains different subpolicy types, a separate Consistency Group will be created for each backup type. Note that if more than one HPE Nimble Storage provider is selected in the job definition, a Consistency Group will be created for each provider. Consistency Groups are named based on the prefix provided during job creation plus the job name.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

10. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page 162.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page 164.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Use the Inventory Browse feature to review the recovery point. See [Browse Inventory](#) on page 357.
- Create a HPE Nimble Storage Restore job definition. See [Create a Restore Job Definition - HPE Nimble Storage](#) on page 306.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**
- [Create a Restore Job Definition - HPE Nimble Storage](#) on page **306**

Create a Backup Job Definition - Pure Storage FlashArray

Back up Pure Storage data with snapshots and replications using a Pure Storage FlashArray Backup job. The RPO and copy data parameters are defined in an SLA Policy, which is then applied to the Backup job definition along with a specified activation time to meet your copy data criteria.

Pure Storage FlashArray can be set to utilize CloudSnap functionality so that snapshots are offloaded from the local storage array to cloud storage. If the Pure Storage SLA was created with the **Add Snapshot Offload**  option, a copy of the snapshot on the local storage array will be offloaded to the cloud and adhere to the number of snapshots or the number of days as provided when creating the sub-policy.

Best Practice: The ECX user interface will indicate that a Pure Storage FlashArray CloudSnap offload job has completed even though the transfer is still occurring in the background which is dependent on network speeds. Consider setting an age as the retention for offload copies when using the Pure Storage FlashArray CloudSnap functionality. Doing so will ensure that a sufficient amount of time has passed for data to be transferred to the S3 storage target before it has to be condensed out from a backup. This is particularly important if several offload jobs are run in quick succession.

BEFORE YOU BEGIN:

- Create and run a Pure Storage FlashArray Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - Pure Storage FlashArray](#) on page **192**.
- Configure an SLA Policy. See [Configure SLA Policies](#) on page **133**.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page **75**.

CONSIDERATIONS:

- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page **155**.

To create a Pure Storage FlashArray Backup job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **Pure Storage FlashArray**.
2. Click **New** , then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more resources to back up.
5. Select an SLA Policy that meets your backup data criteria.
6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs**  tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Note: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum Concurrent Tasks

Set the maximum amount of concurrent transfers between the source and the destination.

Create Consistency Group

If multiple volumes are selected in the Source tab (for example, volumes that contain data tied to an application) enable this option to add the volumes to a Consistency Group to perform backup functions on the entire group. If the associated SLA Policy contains different subpolicy types, a separate Consistency Group will be created for each backup type. Note that if more than one Pure Storage FlashArray provider is selected in the job definition, a Consistency Group will be created for each provider. Consistency Groups are named based on the prefix provided during job creation plus the job name.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot

postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page **152**.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page **334**.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page **553**.

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

10. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** .

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

- Use the Inventory Browse feature to review the recovery point. See [Browse Inventory](#) on page **357**.
- Create a Pure Storage FlashArray Restore job definition. See [Create a Restore Job Definition - Pure Storage FlashArray](#) on page **312**.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**
- [Create a Restore Job Definition - Pure Storage FlashArray](#) on page **312**

Create a Backup Job Definition - VMware

Back up VMware data including virtual machines, datastores, folders, vApps, and datacenters with snapshots using a VMware Backup job. A VMware Backup job consists of snapshot, mirror, and vault sub-policies defined in an SLA Policy, each with their own set of options to give you more control of your VMware protection needs. After an initial primary snapshot is added to the workflow, additional vaults and mirrors ensure your data is replicated to multiple locations.

VMware Backups jobs support IBM, DellEMC Unity, NetApp ONTAP, and Pure Storage FlashArray SLA Policies.

BEFORE YOU BEGIN:

- Create and run a VMware Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - VMware](#) on page **194**.
- Configure an SLA Policy. See [Configure SLA Policies](#) on page **133**.
- Ensure the latest version of VMware Tools is installed in your environment. IBM Spectrum Copy Data Management was tested against VMware Tools 9.10.0.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page **75**.

CONSIDERATIONS:

- Note that VMware Backup and Restore jobs only support vCenters or ESX hosts running vSphere 5.1 through 6.5.
- When running VADP-based VM Replication workflows, target volumes and datastores can be automatically expanded in response to space usage requirements if supported by the underlying storage. Automatic growing prevents a volume from running out of space or forcing you to delete files manually. For a list of supported storage systems see [System Requirements](#) on page **23**.
- Note that VMware DRS cluster datastores are supported in VMware Backup and Restore jobs.
- In NetApp ONTAP environments running Clustered Data ONTAP, cluster peering must be enabled. Peer relationships enable communication between SVMs. See NetApp ONTAP's *Cluster and Vserver Peering Express Guide*.
- In addition to NFS, IBM Spectrum Copy Data Management supports VMFS datastores for NetApp storage targets.
- In IBM storage environments, port grouping and IP partnerships are required to enable remote copy connections. See IBM's [SAN Volume Controller and Storwize Family Native IP Replication Guide](#).

- All related NetApp ONTAP storage resources associated with a VMware provider must be added to IBM Spectrum Copy Data Management, which include NetApp ONTAP storage controllers and clusters. See [Register a Provider](#) on page **75**.
- Note that VMware Backup jobs do not support virtual machine SCSI controllers where the SCSI Bus Sharing value is set to virtual or physical.
- Note that Instant Disk Restore recoveries utilizing the VM Replication method are not supported at the datastore level. Instant Disk Restore datastore level recoveries are supported through the primary storage snapshot method.
- Note that snapshot protection is not supported at an ESX server level.
- Note that cloned volumes will not be replicated through a backup job.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page **155**.

CONSIDERATIONS FOR VMWARE VIRTUAL VOLUMES:

- SLA Policies that include virtual machines stored on virtual volume (VVOL) datastores through VM Replication sub-policies are supported. Replication is supported on the VM Replication target.
- Storage snapshots of virtual machines that reside on a VVOL are currently not supported. If a storage snapshot operation is selected for a virtual machine that resides on a VVOL, the virtual machine is skipped.

AUTOGROW FOR VM REPLICATION (VADP) REQUIREMENTS:

- When running VADP-based VM Replication-based SLA Policies, target volumes and datastores can be automatically expanded in response to space usage requirements if supported by the underlying storage. The following storage systems are supported:
 - NetApp ONTAP Clustered Data ONTAP 8.2 and later (including 7-mode) are supported for all SLA Policy types. Versions earlier than 8.2 are supported for all SLA Policy types except Replication.
 - IBM storage systems running IBM Spectrum™ Accelerate and Virtualize software are supported for all SLA Policy types except Replication.
 - Pure Storage systems are supported for all SLA Policy types.

To create a VMware Backup job replication:

1. Click the **Jobs**  tab. Expand the **Hypervisor** folder, then select **VMware**.
2. Click **New** , then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.

4. From the drop-down menu select **VMs and Templates** or **Storage**. From the list of available sites, select one or more resources to back up, including virtual machines, VM templates, datastores, folders, vApps, and datacenters.
5. Select an SLA Policy that meets your backup data criteria.
6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs**  tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Note: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum Concurrent Tasks

Set the maximum amount of concurrent transfers between the source and the destination.

Create VM snapshots for all VMs

Enable to configure virtual machine snapshot options. Available options include creating virtual machine snapshots for all virtual machines, making all virtual machines included in the job application or file system consistent, or making specific virtual machines included in the job application or file system consistent.

Application consistent backup data captures data in memory and transactions in process. All VSS-compliant applications such as Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, Microsoft SQL, and system state are quiesced. VMDKs and virtual machines can be instantly mounted to recover data related to quiesced applications.

Truncate application logs

To truncate application logs for SQL during the Backup job, enable the **Truncate application logs** option. Note that credentials must be established for the associated virtual machine and SQL instance through the **Sites & Providers**  pane on the **Configure**  tab. Select a VMware provider, click the **VMs** tab, then click the associated virtual machine. Click the **Credentials** tab and add credentials for the virtual machine. Note that System credentials are always required. If the credentials are the same for the SQL instance, select the **Use System Credentials for app** option. If the credentials differ, you must provide credentials for all SQL instances, including the default SQL server. Ensure the Type field in the New Credential dialog window is set to **SQL**.

IBM Spectrum Copy Data Management generates logs pertaining to the application log truncation function and copies them to the following location on the IBM Spectrum Copy Data Management appliance: `/data/log/ecxdeployer/<vm name>/logs`.

VM Snapshot Scripts

VM snapshot prescripts and postscripts are scripts that can be run on the virtual machine before or after a VMware virtual machine snapshot is taken. The snapshot prescript runs before a VMware virtual machine snapshot is captured, while the snapshot postscript runs after the snapshot completes. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines. See [Configure Scripts](#) on page 152.

Select a virtual machine, then click the **Scripts** field in the Pre-Script or Post-Script section to select or upload a script. Once complete, the script displays in the Selected Script(s) section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**.

Click the **Identity** field to add or create the credentials required to run the script. See [Identities Overview](#) on page 127.

Repeat the procedure for each virtual machine associated with the job.

Note: When the VM Snapshot script is run, the virtual machine name will be passed as the first argument to the script. Any additional arguments specified in the job will follow as second, third, and so forth. If a non-zero exit code is returned by the script, the associated snapshot task fails.

Skip read only datastores

Enable to skip datastores mounted as read-only in vCenter.

Skip IA Mount points and/or databases

Enable to skip Instant Disk Restore objects. By default, this option is enabled.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#) on page 334.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

10. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If in a Linux environment, consider creating VADP proxies to enable load sharing. See [Create VMware Backup Job Proxies](#) on page **248**.
- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Use the Inventory Browse feature to review the recovery point. See [Browse Inventory](#) on page **357**.
- Create a VMware Restore job definition. See [Create a Restore Job Definition - VMware](#) on page **317**.

RELATED TOPICS:

- [NetApp ONTAP Document: Cluster and Vserver Peering Express Guide](#)
- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Configure Scripts](#) on page **152**
- [Create a Schedule](#) on page **155**
- [Create VMware Backup Job Proxies](#) on page **248**
- [Create a Restore Job Definition - VMware](#) on page **317**

Create VMware Backup Job Proxies

In IBM Spectrum Copy Data Management, running VMware Backup jobs through VADP can be taxing on your system resources. By creating VMware Backup job proxies, you enable load sharing and load balancing for those jobs in Linux environments.

Note that the first time a given job is run, the proxies do not take effect because VM clone technology is used. But the second and subsequent times the job is run, change block tracking technology is used and the proxies are employed.

If proxies exist, the entire processing load is shifted off the IBM Spectrum Copy Data Management host machine and onto the proxies, else the entire load stays on the IBM Spectrum Copy Data Management host. Within a Backup job, the processing load for any single VM is shifted to a single proxy machine; multiple VMs are shifted to multiple proxies if available.

If a proxy server goes down or is otherwise disabled before the start of the job, the other proxies (or if there are no other proxies, the IBM Spectrum Copy Data Management host) take over and the job completes. If a proxy server becomes disabled during the running of a job, there is a possibility that the job will fail.

BEFORE YOU BEGIN:

- Determine how many proxies to create; the more proxies, the faster the jobs run. Each proxy is used merely to process data, and the results are returned through the host IBM Spectrum Copy Data Management server. The only impact seen by the user is the improvement in performance when running the job.

System Requirements:

This feature has been tested only for Ubuntu, SUSE Linux Enterprise Server, and Red Hat environments. It is supported only in x64 configurations with a minimum kernel of 2.6.32.

A minimum of 8 GB of RAM is required (16 GB recommended), along with 60 GB of disk space.

Each proxy must have a fully qualified domain name.

Installer Notes:

The IBM Spectrum Copy Data Management 2.2.6+ version of the VADP Proxy installer includes Virtual Disk Development Kit (VDDK) version 5.5.5. This version of the VADP proxy installer provides the following functionality:

- External VADP Proxy support with vSphere 6.5
- External VADP Proxy support for Hot Add operations, which provides higher performance for VADP Backups.

To create a proxy:

For each proxy:

1. Power up a physical or virtual Linux machine that meets the system requirements defined above, and is on the same network as the host IBM Spectrum Copy Data Management machine.
2. Copy the VADP Proxy installation program to the local proxy machine. This can be attained from the IBM Spectrum Copy Data Management Marketplace.
3. Log in to the proxy machine as root, or as a user capable of running “sudo” commands.
4. On the proxy machine, open a terminal. Enter the following command to install the proxy server software:

```
./vmdkbackup-1.0-installer.bin
```

The Setup wizard opens.

Note: Alternatively you can run the installer using command line protocol by entering the following command: `./vmdkbackup-1.0-installer.bin --mode text`

5. Follow the steps in the Setup wizard to configure your proxy server and connect to the IBM Spectrum Copy Data Management host.
 - a. When prompted for the installation directory, select **/opt/CDM**.
 - b. When prompted for the IBM Spectrum Copy Data Management Discovery Server IP, enter the IP Address of the IBM Spectrum Copy Data Management host.
 - c. When prompted for the IBM Spectrum Copy Data Management Site String, set it to “default”.
 - d. When prompted for Performance Optimization, set it to “unknown”.
6. Click **Finish** when the Setup wizard indicates it has completed. After installation, note that your new installation directory includes a subdirectory called `/log`, which is the job log location.

After successful installation, the service **ecxvadp** is started on the proxy machine. A log file **ecxvadp.log** is generated in `/opt/CDM/logs` directory.

Repeat the previous steps for each proxy you want to create.

To confirm proxy server connections:

1. Start the IBM Spectrum Copy Data Management management console.
2. In the IBM Spectrum Copy Data Management management console, click the arrow next to the **Support**  icon, then choose **View Edge Services Status**. The IBM Spectrum Copy Data Management host machine and each operational proxy server are displayed.

NEXT STEPS:

- Run the VMware Backup job. The use of the proxies are indicated in the job log by a log message similar to the following:

```
Run remote vmdkbackup of MicroService: http://<proxy
node name>, IP:<proxy IP address>
```

- Uninstall the proxies when you cease running the VMware Backup jobs. To uninstall a proxy, on your host machine, run the following command from the uninstall subdirectory of the installation directory (default installation directory is /opt/CDM):

```
./uninstall_vmdkbackup
```

The installation directory is removed.

RELATED TOPICS:

- [Create a Backup Job Definition - VMware](#) on page **242**
- [Start, Pause, and Hold a Job Session](#) on page **162**

Restore Jobs

The topics in the following section cover Restore job definitions as well as postscript argument details.

Create a Restore Job Definition - InterSystems Caché

IBM Spectrum Copy Data Management leverages Copy Data Management technology for recovering application databases through Database Restore jobs. Your InterSystems Caché clones can be utilized and consumed instantly through IBM Spectrum Copy Data Management Instant Disk Restore jobs. IBM Spectrum Copy Data Management catalogs and tracks all cloned instances. Instant Access leverages iSCSI or fibre channel protocols to provide immediate mount of LUNs without transferring data. Snapshotted databases are cataloged and instantly recoverable with no physical transfer of data.

BEFORE YOU BEGIN:

- Create and run an InterSystems Caché Backup job. See [Create a Backup Job Definition - InterSystems Caché](#) on page **197**.
- Review InterSystems Caché requirements. See [InterSystems Caché Requirements](#) on page **34**.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page **75**.
- You must add credentials to the destination virtual machine when recovering with the subnet option. See [Add Credentials to a Virtual Machine](#) on page **101**.

CONSIDERATIONS:

- Note that the following users and groups must be created on the target host: instance owner, effective user for InterSystems Caché superserver and its jobs, effective group for InterSystems Caché processes, and a group that has permissions to start and stop InterSystems Caché instances. The user and group IDs should match those on the source host. The instance will be brought up using the same mount points as those found on the source machine, so ensure these mounts are not in use on the target.
- Note that it is possible to scan in an InterSystems Caché backup failover member instance or an async member instance and run snapshots against the mirror copy instead of the primary failover member.
- When creating an InterSystems Caché restore job definition, select only one instance to restore. If more than once instance is selected, the InterSystems Caché agent only restores the last instance it receives in the command request.
- When restoring to a target with running InterSystems Caché instances, the instances display as valid targets. Note that IBM Spectrum Copy Data Management will not interact with these instances, but instead bring up a new instance using mapped mount points. When restoring to a target with no prior InterSystems Caché instances, IBM Spectrum Copy Data Management creates a placeholder that acts as a restore target named

cache_general. Note that *cache_general* should only be used as a restore target and should not be selected for backup.

- Single InterSystems Caché databases can be restored through an Instant Disk Restore job, which mounts physical volumes on the target machine. Granular recovery can then be performed through InterSystems Caché commands.

Best Practice: Create a schedule before creating a job definition so that you can easily add the schedule to the job definition.

To create an InterSystems Caché Restore job definition:

1. Click the **Jobs**  tab. Expand the **Database** folder, then select **InterSystems Caché**.
2. Click **New** , then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select a template. Available options include **Instant Database Restore**  and **Instant Disk Restore** .
5. Click **Source** . From the drop-down menu select **Application Browse** to select a source site and an application server to view available database recovery points. Select resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.
Alternatively, select **Application Search** from the drop-down menu to search for application servers with available recovery points. Add copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.
6. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Destination** . Select a source site and an associated destination. If creating an **Instant Disk Restore**  job definition, review the destination's database name mapping settings. Optionally, click the **New database name** field to create an alternate database name.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Application Options

Rename Mount Points

For more information about the **Rename Mount Points** options, see [Restore Jobs - Rename Mount Points and Initialization Parameter Options](#) on page 331.

If creating an Instant Database Restore job definition, this option is set to Do Not Rename by default. IBM Spectrum Copy Data Management will mount the mount points with the same path/name as the source.

Change the GUID of the Instance

By default the restored instance will have the same GUID as the source instance. If creating an Instant Database Restore job definition, selecting this option will generate a new GUID to be assigned to the restored instance.

Policy Options

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If unselected, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the database recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Allow to overwrite vDisk

In cases where the Make Permanent option is enabled, and the destination VM has conflicting VMDK files, enable the **Allow to overwrite vDisk** option to delete the existing VMDK and overwrite it with the selected source.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334.

State and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

Storage Options

Make Permanent

Set the default permanent restoration action of the job. All database recovery operations can leverage Instant or Test modes and then either be deleted or promoted to permanent mode. This behavior is controlled through the **Make Permanent** option.

Enabled - Always make permanent through full copy FlashCopy

Disabled - Never make permanent

User Selection - Allows the user to select Make Permanent or Cleanup when the job session is pending

Protocol Priority

If more than one storage network protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

- Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

- Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

- When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page 162.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page 164.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Create a Backup Job Definition - InterSystems Caché](#) on page **197**
- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**

Create a Restore Job Definition - SAP HANA

IBM Spectrum Copy Data Management leverages Copy Data Management technology for recovering application databases through Database Restore jobs. Your SAP HANA clones can be utilized and consumed instantly through IBM Spectrum Copy Data Management Instant Disk Restore jobs. IBM Spectrum Copy Data Management catalogs and tracks all cloned instances. Instant Disk Restore leverages ISCSI or fibre channel protocols to provide immediate mount of LUNs without transferring data. Snapshotted databases are cataloged and instantly recoverable with no physical transfer of data.

BEFORE YOU BEGIN:

- Create and run a SAP HANA Backup job. See [Create a Backup Job Definition - SAP HANA](#) on page 200.
- Review SAP HANA requirements. See [SAP HANA Requirements](#) on page 51.
- Review SAP HANA prerequisites. See [Prerequisites](#) on page 52.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.
- You must add credentials to the destination virtual machine when recovering with the subnet option. See [Add Credentials to a Virtual Machine](#) on page 101.
- When performing a SAP HANA Disk Restore, point-in-time (PIT) is an option. For restore jobs of SAP HANA log backups, PIT cannot be used. Instead, restore at the snapshot level and use SAP HANA Studio to do a PIT restore.

Best Practice: Create a schedule before creating a job definition so that you can easily add the schedule to the job definition.

To create a SAP HANA Restore job definition:

1. Click the **Jobs**  tab. Expand the **Database** folder, then select **SAP HANA**.
2. Click **New** , then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select a template. Available options include **Instant Database Restore**  and **Instant Disk Restore** .
5. Click **Source** . From the drop-down menu select **Application Browse** to select a source site and an application server to view available database recovery points. Select resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **Application Search** from the drop-down menu to search for application servers with available recovery points. Add copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Destination** . Select a source site and an associated database. Review the destination's database name mapping settings.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Application Options

Rename Mount Points

For more information about the **Rename Mount Points** options, see [Restore Jobs - Rename Mount Points and Initialization Parameter Options](#) on page 331.

Policy Options

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If unselected, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the database recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Allow to overwrite vDisk

In cases where the Make Permanent option is enabled, and the destination VM has conflicting VMDK files, enable the **Allow to overwrite vDisk** option to delete the existing VMDK and overwrite it with the selected source.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page **553**.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page **334**. **State** and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

Storage Options

Make Permanent

Set the default permanent restoration action of the job. All database recovery operations can leverage Instant or Test modes and then either be deleted or promoted to permanent mode. This behavior is controlled through the **Make Permanent** option.

Enabled - Always make permanent through full copy FlashCopy

Disabled - Never make permanent

User Selection - Allows the user to select Make Permanent or Cleanup when the job session is pending

Protocol Priority

If more than one storage network protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

- Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

- Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

- When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

Restoring a SAP HANA instance through the SAP HANA Client (Studio)

After the IBM Spectrum Copy Data Management SAP HANA Instant Disk Restore job completes, you must restore the data volume, then finalize the restore using SAP HANA Client. Note that this procedure applies only to Instant Disk Restore jobs utilizing the Make Permanent option.

Note: Users of SAP HANA 2.0 SP1 or later must register their systems with the Multiple Containers mode and the System database option selected. Further options specific to SAP HANA 2.0 SP1 or later are specified in the procedure below.

- On the SAP HANA Client, right-click the SAP HANA system and select **Backup and Recovery**, then select **Recover System**. SAP HANA 2.0 SP1 users should select **Recover System Database**. The SAP HANA system shuts down in preparation for the recovery.
- From a command line, perform the following steps:
 - Unmount the original data volume (if present) through the following command: `umount /<path of data installation>` e.g. `umount /hana/data`
 - Unmount the data volume recovered by IBM Spectrum Copy Data Management through the following command: `umount /<path of recovered volume>` e.g. `umount /dev/sde1`
 - Mount the new data volume again through the following command: `mount /dev/mapper/newID /<path of original data installation>` e.g. `mount /dev/sde1 /hana/data`
- For SAP HANA 2.0 SP1 or later environments:** Back in the SAP HANA Client interface, keep the default selections and click **Next** until the last of the backup options display, then click **Finish**.

Restoring an SAP HANA system database from a specific backup or storage snapshot:

- Right-click the SAP HANA system and select **Backup and Recovery**.
- Select **Recover System Database** to continue with the recovery. The Specify system database dialog will load.
- Select the SAP HANA system database.
- Click **Next**. The Specify Recovery Type dialog will appear.
- Select the **Recover the database to a specific data backup or storage snapshot** option, then click **Next**.
- In the **Backup Location** screen, keep the default selections and click **Next**.
- In the **Select a Backup** screen, select the backup marked with a green availability then continue clicking **Next** until the last of the backup options displays, then click **Finish**.

Restoring an SAP HANA tenant database from a specific backup or storage snapshot:

1. Right-click the SAP HANA system and select **Backup and Recovery**.
2. Select **Recover Tenant Database** to continue with the recovery. The Specify tenant database dialog will load.
3. Select the SAP HANA tenant database system.
4. Click **Next**. The Specify Recovery Type dialog will appear.
5. Select the **Recover the database to a specific data backup or storage snapshot** option, then click **Next**.
6. In the **Backup Location** screen, keep the default selections and click **Next**.
7. In the **Select a Backup** screen, select the backup marked with a green availability then continue clicking **Next** until the last of the backup options displays, then click **Finish**.

Recovering an SAP HANA system database to a point-in-time (PIT):

1. Right-click on the SAP HANA system and select **Backup and Recovery**.
2. Select **Recover System Database**. The Recovery of SYSTEMDB@<SID> window will open to the Specify Recovery Type dialog. Select the SAP HANA system database.
3. Click **Next**. The Specify Recovery Type dialog will appear.
4. Select **Recover the database to the following point in time**. Enter a **Date**, a **Time**, and select the time zone from the **Select Time Zone** selection field.

Note: The date and time selection should be within the defined range of when the point in time was started.

5. Click **Next**. The Locate Backup Catalog dialog will appear.
6. Select **Recover using the backup catalog** and **Search for the backup catalog in the file system only**.
7. For **Backup Catalog Location**, provide the catalog location. For example:

```
/hana/logbackup/<SID>/catalog/SYSTEMDB
```

where <SID> is the system ID of the database. This is the directory that is created on the SAP HANA system in the prerequisites. For more information, see [Prerequisites](#) on page 52.
8. Click **Next**. The Select a Backup dialog will appear.
9. Select the desired backup based on the **Start Time** that is marked as **Available**. If you do not see the backup listed, click the **Refresh** button or the **Show More** button if several backups are listed.
10. Click **Next**. The Locate Log Backups dialog will appear.
11. In the **Locations** field, enter the log backup location. For example:

```
/hana/logbackup/<SID>/SYSTEMDB
```

where <SID> is the system ID of the database.
12. Click **Add**. It may be necessary to select this location from the list of locations.

13. Click **Next**. The Other Settings dialog will appear. Ensure **File System** and **Use Delta Backups (Recommended)** is selected.
14. Click **Next**. The Review Recovery Settings dialog will appear. Review the selections and click **Back** if any revisions are necessary.
15. Click **Finish**. The Recovery Execution Summary will appear.
16. Click **Close**.

Restoring an SAP HANA tenant database to a point-in-time (PIT):

1. Right-click on the SAP HANA system and select **Backup and Recovery**.
2. Select **Recover tenant database**. The Specify tenant database dialog will load. Select the SAP HANA tenant database system.
3. Click **Next**. The Specify Recovery Type dialog will appear.
4. Select **Recover the database to the following point in time**. Enter a **Date**, a **Time**, and select the time zone from the **Select Time Zone** selection field.
Note: The date and time selection should be within the defined range of when the point in time was started.
5. Click **Next**. The Locate Backup Catalog dialog will appear.
6. Select **Recover using the backup catalog** and **Search for the backup catalog in the file system only**.
7. For **Backup Catalog Location**, provide the catalog location. For example:
`/hana/logbackup/<SID>/catalog/DB_<SID>`
where *<SID>* is the system ID of the database. This is the directory that is created on the SAP HANA system in the prerequisites. For more information, see [Prerequisites](#) on page 52.
8. Click **Next**. The Select a Backup dialog will appear.
9. Select the desired backup based on the **Start Time** that is marked as **Available**. If you do not see the backup listed, click the **Refresh** button or the **Show More** button if several backups are listed.
10. Click **Next**. The Locate Log Backups dialog will appear.
11. In the **Locations** field, enter the log backup location. For example:
`/hana/logbackup/<SID>/DB_<SID>`
where *<SID>* is the system ID of the database.
12. Click **Add**. It may be necessary to select this location from the list of locations.
13. Click **Next**. The Other Settings dialog will appear. Ensure **File System** and **Use Delta Backups (Recommended)** is selected.
14. Click **Next**. The Review Recovery Settings dialog will appear. Review the selections and click **Back** if any revisions are necessary.

15. Click **Finish**. The Recovery Execution Summary will appear.
16. Click **Close**.

Tip: After completing an Instant Disk Restore for SAP HANA and using the Make Permanent option, both `data` and `temp` data disks may still appear on the SAP HANA server until the SCSI bus is rescanned or the SAP HANA server is restarted. Rescan the SCSI bus on the SAP HANA server:

```
# ls /sys/class/scsi_host/host
For each host found, where <x> is the number of the host, issue:
# echo "- - -" > /sys/class/scsi_host/host<x>/scan
```

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Create a Backup Job Definition - SAP HANA](#) on page **200**
- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**

Create a Restore Job Definition - Oracle

IBM Spectrum Copy Data Management leverages Copy Data Management technology for recovering application databases through Database Restore jobs. Your Oracle clones can be utilized and consumed instantly through IBM Spectrum Copy Data Management Instant Disk Restore jobs. IBM Spectrum Copy Data Management catalogs and tracks all cloned instances. Instant Disk Restore leverages ISCSI or fibre channel protocols to provide immediate mount of LUNs without transferring data. Snapshotted databases are cataloged and instantly recoverable with no physical transfer of data. Point-in-time recovery is supported with log forwarding.

The following Oracle Database workflows are supported: **DevOps**, which provides Instant Disk Restore or Instant Recovery to a new location using a masked image, created through a Restore job with data masking enabled, **Instant Database Recovery**, which provides Instant Disk Restore or Instant Recovery using a non-masked image and point-in-time transaction logs, and **Instant Disk Restore**, which mounts a database for RMAN restores with application and operating system support.

BEFORE YOU BEGIN:

- Create and run an Oracle Backup job. See [Create a Backup Job Definition - Oracle](#) on page **204**.
- Review Oracle requirements. See [Oracle Requirements](#) on page **41** and [Oracle Database Support FAQ](#) on page **563**.
- To ensure that filesystem permissions are retained correctly when IBM Spectrum Copy Data Management moves Oracle data between servers, ensure that the user and group IDs of the Oracle users (e.g. oracle, oinstall, dba) are consistent across all the servers. Refer to Oracle documentation for recommended uid and gid values.
- If Oracle data resides on LVM volumes, you must stop and disable the lvm2-lvmetad service before running Backup or Restore jobs. Leaving the service enabled can prevent volume groups from being resignatured correctly during restore and can lead to data corruption if the original volume group is also present on the same system. To disable the lvm2-lvmetad service, run the following commands:

```
systemctl stop lvm2-lvmetad
systemctl disable lvm2-lvmetad
```

Next, disable lvmetad in the LVM config file. Edit the file /etc/lvm/lvm.conf and set:

```
use_lvmetad = 0
```

- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page **75**.
- You must add credentials to the destination virtual machine when recovering with the subnet option. See [Add Credentials to a Virtual Machine](#) on page **101**.

ORACLE DATABASE CONSIDERATIONS:

- For Oracle 12c databases, backups are created without placing the database in hot backup mode through Oracle Storage Snapshot Optimization. All associated snapshot functionality is supported. This feature requires the *Advanced Compression* feature of Oracle to be licensed. If this feature is not licensed in your environment, perform the following procedure to disable Snapshot Optimization and force the use of hot backup mode:

Create the file `/etc/guestapps.conf` on the Oracle server and add the following to it:

```
[DEFAULT]

skipSnapshotOptimization = true
```

If the file already exists, edit it and add the parameter under the existing [DEFAULT] section. This is a per-host setting. The parameter must be set in this file on each Oracle server where you want to force the use of hot backup mode.

- Note that point-in-time recovery is not supported when one or more datafiles are added to the database in the period between the chosen point-in-time and the time that the preceding Backup job ran.
- To properly recover Oracle databases, the ORACLE_BASE environment variable must be set in the login environment of the user that owns the Oracle home. In cases where multiple homes exist on a node, each with a different owner, the variable must be set for each of the owners.
- During an Instant Database Restore, there may be failures if the new name specified for the restored database is similar to an existing database only differing by numerical suffix. For clustered instances of Oracle databases, the appliance always uses global database name in the UI. During the inventory and restore processes, individual instances using the numerical suffixes of the cluster must be correlated to the global database name. The issue with this comes when, as an example, "Production12" is discovered. Is this the instance 12 of the "Production" database, or instance two (2) of the "Production1" database, or a database named "Production12."

Best Practice: Create a schedule before creating a job definition so that you can easily add the schedule to the job definition.

To create an Oracle Restore job definition:

1. Click the **Jobs**  tab. Expand the **Database** folder, then select **Oracle**.
2. Click **New** , then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.

4. Select a template. Available options include **DevOps** , **Instant Database Restore** , and **Instant Disk Restore** .
5. Click **Source** . From the drop-down menu select **Application Browse** to select a source site and an application server to view available database recovery points. Select resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **Application Search** from the drop-down menu to search for application servers with available recovery points. Add copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.
6. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.

If creating an **Instant Database Restore**  job definition, an additional recovery option is available through the Select Version feature. Enable **Allow Point-in-Time selection when job runs** to leverage archived logs and enable a point-in-time recovery of the databases.

If creating an **Instant Disk Restore**  job definition, the RMAN tag displays next to the time in the Version field. An Oracle administrator can correlate the RMAN backups to the IBM Spectrum Copy Data Management versions during job creation.
7. Click **Destination** . Select a source site and an associated Oracle home. Click the Destination field to enter an optional alternate name for the database.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Application Options

Record mounted copies in RMAN local repository

This option is available for Instant Disk Restore workflows.

Select this option to catalog mounted copies into RMAN at the end of the Instant Disk Restore job. This option is an alternative to the cataloging of IBM Spectrum Copy Data Management-created copies during a Backup job. If the RMAN cataloging option is selected in the Backup job definition, every copy created by IBM Spectrum Copy Data Management is cataloged in the source database immediately after the copy is created. By contrast, this option allows you to perform cataloging on-demand only for a specific copy and only when you intend to restore data from that copy through RMAN. Note that for the cataloging to succeed the target database must be running at the time the Instant Disk Restore job runs.

Rename Mount Points and Database Initialization Parameters

For more information about the **Rename Mount Points** and **Database Initialization Parameters** (Instant Database Recovery and DevOps workflows only) options, see [Restore Jobs - Rename Mount Points and Initialization Parameter Options](#) on page 331.

ASM Disk Names

This option allows you to specify the disk naming pattern for restored ASM disks, if available. If **Use default pattern** is selected, IBM Spectrum Copy Data Management uses the default naming pattern, which is `/dev/ecx-asmdisk/*` in Linux environments or `/dev/ecx_asm*` in AIX environments. Select **Specify a custom pattern** to set ASM disks to follow any naming conventions that may be in use for existing disks. The custom pattern must begin with `"/dev"` and must end with an asterisk (`*`). During restore, IBM Spectrum Copy Data Management creates a device alias, or symlink, matching the specified pattern and replaces the asterisk with a unique disk name.

Note: This option has no effect if the database being restored does not use any ASM disks.

Policy Options

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If unselected, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the database recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Replace existing database

This option is available for Instant Database Recovery workflows.

Select this option to replace an existing database with the same name during recovery. When an Instant Database Recovery is performed for a database and another database with the same name is already running on the destination host/cluster, IBM Spectrum Copy Data Management shuts down the existing database before starting up the recovered database. If this option is not selected, the Instant Database Recovery fails when IBM Spectrum Copy Data Management encounters an existing running database with the same name.

Leave database shut down after recovery

This option is available for Instant Database Recovery and DevOps workflows.

Select this option to shut down the recovered database once the recovery operation completes. The database can be started up manually once needed.

Allow to overwrite vDisk

In cases where the Make Permanent option is enabled, and the destination VM has conflicting VMDK files, enable the **Allow to overwrite vDisk** option to delete the existing VMDK and overwrite it with the selected source.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334. **State** and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

Storage Options

Make Permanent

Set the default permanent restoration action of the job. All database recovery operations can leverage Instant or Test modes and then either be deleted or promoted to permanent mode. This behavior is controlled through the **Make Permanent** option.

Enabled - Always make permanent through full copy FlashCopy

Disabled - Never make permanent

User Selection - Allows the user to select Make Permanent or Cleanup when the job session is pending

Protocol Priority

If more than one storage network protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Create a Backup Job Definition - Oracle](#) on page **204**
- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**

Create a Restore Job Definition - SQL

IBM Spectrum Copy Data Management leverages Copy Data Management technology for recovering application databases through Database Restore jobs. Your SQL clones can be utilized and consumed instantly through IBM Spectrum Copy Data Management Instant Disk Restore jobs. IBM Spectrum Copy Data Management catalogs and tracks all cloned instances. Instant Disk Restore leverages iSCSI or fibre channel protocols to provide immediate mount of LUNs without transferring data. Snapshotted databases are cataloged and instantly recoverable with no physical transfer of data. Point-in-time recovery is supported with log forwarding.

The following Microsoft SQL Server workflows are supported:

Instant Database Restore - Provides instant recovery using a non-masked image and point-in-time transaction logs.

When performing a restore to a primary replica of a specific AlwaysOn Availability Group, the restore database is added to the destination availability group, however, the secondary replica is not automatically replicated.

When performing a restore to a secondary replica, the database is restored with the "norecovery" option, and the database is set to a "restoring" state. If managing transaction log backups without using IBM Spectrum Copy Data Management, you can manually restore log files, and add the database to an availability group, assuming the Isn of the secondary and primary database copies meet the criteria.

Instant Disk Restore - Mounts a database for restores with application and operating system support. An Instant Disk Restore of an AlwaysOn database is restored to the local destination instance.

Instant Seeding - Instant Seeding is used when a primary database is up and running, a secondary replica exists (but no secondary database). The primary database can be protected through an IBM Spectrum Copy Data Management Copy job, with log backups performed regularly. The seeding process restores the database to the copy time and applies the transaction logs of the primary database. The database is then added to an availability group at the end of the restore.

Instant Seeding is a recovery template available for SQL AlwaysOn Availability Groups that initializes secondary replica instances using storage copies. During an Instant Seeding restore, snapshot data along with log backups are restored to the secondary database, then added to the AlwaysOn Availability Group. Other AlwaysOn Availability Group restore jobs place databases in "restoring" states or only create local instances. Additional processes at the end of the Instant Seeding restore job add the database of a secondary replica to the Availability Group.

When creating the Instant Seeding job definition, primary and secondary replicas of the original Availability Group are displayed as valid restore destinations. All log backups after the copy job time are restored; no point-in-time options are required. For more information, see the *Instant Seeding Prerequisites and Limitations* section of the [Microsoft SQL Server Requirements](#) on page 57.

BEFORE YOU BEGIN:

- Create and run a SQL Backup job. See [Create a Backup Job Definition - SQL](#) on page 211.

- Review SQL requirements. See [Microsoft SQL Server Requirements](#) on page **57** and [Microsoft SQL Server Support FAQ](#) on page **584**.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page **75**.
- You must add credentials to the destination virtual machine when recovering with the subnet option. See [Add Credentials to a Virtual Machine](#) on page **101**.

MICROSOFT SQL SERVER CONSIDERATIONS:

- You must use the **Overwrite existing database** option to restore a Microsoft SQL database to a source instance. Note that in an Instant Disk Restore or Instant Database Restore job, the restore destination cannot be the source instance if this option is not selected. Due to Microsoft SQL restrictions, two databases with the same ID cannot exist on the same instance, including if the restored database is renamed through the job definition.
- You must use a proxy node when performing a restore to a SQL Failover cluster. Windows requires signatures to be unique. When a disk is attached that has a signature equal to one that is already attached, Windows keeps the disk in “offline” mode and does not read its partition table or mount its volumes. To prevent the issue of disk signature collision, during Instant Database Restore, IBM Spectrum Copy Data Management leverages a Windows proxy server to temporarily mount disks from snapshots, generate a new signature, then mount to original server.
- You must manually create a directory to store VSS provider logs when running IBM Spectrum Copy Data Management 2.2.6 and earlier. Create the following directory structure on the SQL server: c:\temp\CDM\logs
- Note that Instant Disk Restores of AlwaysOn databases are restored to local instances only. Instant Database Recoveries of AlwaysOn databases are joined to destination Availability Groups if the restore target is a primary replica, and restored to the local instance in a "RESTORING" state if the restore target is a secondary replica.
- When running SQL Server 2012, certain Availability Group properties must be modified in order for log backups to generate in their specified location. Under **Backup Preferences**, set the **Where should backups occur?** option to **Any Replica**. Under **Replica backup priorities**, set the primary server instance to 55 and the secondary server instance to 45.
- It is recommended to keep the SQL transaction log location on a different disk than the SQL database data and log files.
- When creating an Instant Seeding restore job definition, the destination must be a non-system drive. The SQL database and log files must be located on non-system drives.

Best Practice: Create a schedule before creating a job definition so that you can easily add the schedule to the job definition.

To create a Microsoft SQL Restore job definition:

1. Click the **Jobs**  tab. Expand the **Database** folder, then select **SQL**.
2. Click **New** , then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select **Microsoft SQL (Standalone and Failover Cluster)**  or **Microsoft SQL (Always On Availability Group)** .
5. Select a template. Available options include **Instant Database Restore**  for Microsoft SQL Standalone/Failover Cluster and Always On Availability Group jobs, **Instant Disk Restore** , and **Instant Seeding**  for Microsoft SQL Always On Availability Group jobs.
6. Click **Source** . From the drop-down menu select **Application Browse** to select a source site and an application server to view available database recovery points. Select resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.
Alternatively, select **Application Search** from the drop-down menu to search for application servers with available recovery points. Add copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.
7. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
If creating an **Instant Database Restore**  job definition, an additional recovery option is available through the Select Version feature. Enable **Allow Point-in-Time selection when job runs** to leverage archived logs and enable a point-in-time recovery of the databases.
8. Click **Destination** . Select a source site and an associated Microsoft SQL database. Click the **New database name** field to enter an optional alternate name for the database. If the destination is a SQL Failover Cluster instance, select a Windows server proxy in the **Select a Windows server to resigature LUNs** section. When resignaturing a copy, IBM Spectrum Copy Data Management retains the data and mounts the volume to the proxy selected in the Select a Windows server to resigature LUNs section.

Note: Any Windows node with iSCSI or Fibre Channel access to the storage can be selected as a proxy server, provided that the node is not part of the original cluster. It is recommended to select a standalone virtual or physical Windows node as a proxy server.

Note: If creating an **Instant Seeding**  job definition, the seeding target must be a secondary replica. If a primary AlwaysOn node is selected as a seeding target, the job will fail. The **New database name** option does not apply to Instant Seeding jobs.

9. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
10. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Application Options

Roll back uncommitted transactions and leave the database ready to use

Select this option to restore the database to an online state. If selected, additional transaction logs cannot be restored. If deselected, uncommitted transactions are not rolled back, leaving the database non-operational. Additional transaction logs can then be restored.

Note: This option does not apply to Instant Seeding jobs.

Overwrite existing database

Select this option to replace an existing database with the same name during recovery. When an Instant Database Recovery is performed for a database and another database with the same name is already running on the destination host/cluster, IBM Spectrum Copy Data Management shuts down the existing database before starting up the recovered database. If this option is not selected, the Instant Database Recovery fails when IBM Spectrum Copy Data Management encounters an existing running database with the same name.

Copy databases

Select this option to restore SQL files to the original file path that is currently in use. The IBM Spectrum Copy Data Management agent copies the database files from the snapshot or clone volumes to the database folders on the local drive.

Rename mount points

By default, IBM Spectrum Copy Data Management renames mount points to the SQL data directory of the target SQL instance. You can override this behavior through the Rename mount points option.

Default SQL data directory - Mount points are not renamed.

Original mount point or drive letter - The original volume mount point of the databases is used. This will keep the same drive mapping, but the database recovery will occur on an alternate server. If the original volume mount point cannot be used to mount a volume (for example, if the folder is not empty), the restore fails. Note that it is recommended to keep the SQL transaction log location on a different disk than the SQL database data and log files.

Add a custom mount point - If enabled, a custom prefix can be entered in the **Prefix string** field. The prefix must specify a valid, preexisting volume drive letter that can contain a volume mount point on the drive. The prefix substitutes the default root folder of the mount. For example, if the default root folder of the mount drive is E:\SQLDataFiles\mnt, and F:\RestoreMnt is entered in the Mount Point Prefix field,

E:\SQLDataFiles\mnt is renamed to F:\RestoreMnt. The drive, in this case the F drive, must exist on the destination server. If the folder, in this case the RestoreMnt folder, does not exist, it will be created.

Note: This option does not apply to Instant Seeding jobs. Instant Seeding supports restoring to the original path and volume drive letter/mount points.

Policy Options

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the database recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Allow to overwrite vDisk

In cases where the Make Permanent option is enabled, and the destination VM has conflicting VMDK files, enable the **Allow to overwrite vDisk** option to delete the existing VMDK and overwrite it with the selected source.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334.

State and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

Storage Options

Make Permanent

Set the default permanent restoration action of the job. All database recovery operations can leverage Instant, or Test, mode and then either deleted or promoted to permanent mode. This behavior is controlled through the **Make Permanent** option.

Enabled - Always make permanent

Disabled - Never make permanent

User Selection - Allows the user to select Make Permanent or Cleanup when the job session is pending

Protocol Priority

If more than one storage networking protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

11. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

12. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

13. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page 162.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page 164.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Create a Backup Job Definition - SQL](#) on page **211**
- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**

Create a Restore Job Definition - File System

IBM Spectrum Copy Data Management leverages Copy Data Management technology for recovering physical Windows, AIX, and Linux file systems through File System Restore jobs. Clones can be utilized and consumed instantly through IBM Spectrum Copy Data Management Instant Disk Restore jobs. IBM Spectrum Copy Data Management catalogs and tracks all cloned instances. Instant Disk Restore jobs mount file systems from storage copies containing application data.

BEFORE YOU BEGIN:

- Review File System requirements. See [File System Requirements](#) on page 28.
- Create and run a File System Backup job. See [Create a Backup Job Definition - File System](#) on page 216.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.
- You must add credentials to the destination virtual machine when recovering with the subnet option. See [Add Credentials to a Virtual Machine](#) on page 101.

To create a File System Restore job definition:

1. Click the **Jobs**  tab. Expand the **File System** folder, then select **File System**.
2. Click **New** , then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Instant Disk Restore**  template.
5. Click **Source** . From the drop-down menu select **Application Browse** to select a source site and a file system to view available recovery points. Select resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.
Alternatively, select **Application Search** from the drop-down menu to search for file systems with available recovery points. Add copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.
6. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Destination** . Select a source site and an associated destination. Review the destination's mount point mapping settings. Optionally, click the **Enter an alternate mount point** field to create an alternate mount point, or select **Use original mount points**.

8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Application Options

Overwrite Existing Mount Points

Select to overwrite the mount points at their original location.

Policy Options

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If unselected, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the database recovery fails.

Allow to overwrite vDisk

In cases where the Make Permanent option is enabled, and the destination VM has conflicting VMDK files, enable the **Allow to overwrite vDisk** option to delete the existing VMDK and overwrite it with the selected source.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334.

State and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

Note: If adding a script to a Windows-based File System job definition, the user running the script must have the "Log on as a service" right enabled, which is required for running prescripts and postscripts. For more information about the "Log on as a service" right, see <https://technet.microsoft.com/en-us/library/cc794944.aspx>.

Storage Options

Make Permanent

Set the default permanent restoration action of the job. All database recovery operations can leverage Instant or Test modes and then either be deleted or promoted to permanent mode. This behavior is controlled through the **Make Permanent** option.

Enabled - Always make permanent through full copy FlashCopy

Disabled - Never make permanent

User Selection - Allows the user to select Make Permanent or Cleanup when the job session is pending

Protocol Priority

If more than one storage network protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

- Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

- Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

- When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page 162.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page 164.

- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Create a Backup Job Definition - File System](#) on page **216**
- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**

Create a Restore Job Definition - DellEMC Unity

IBM Spectrum Copy Data Management leverages Copy Data Management technology for recovering DellEMC Unity volumes through DellEMC Unity Restore jobs.

Instant Disk Restore

Provides instant writable access to a volume. An IBM Spectrum Copy Data Management snapshot is mapped to a target server where it can be accessed, copied, or put immediately into production use as needed.

Restore Volume(s)

Recover a volume from a snapshot or replication created through an IBM Spectrum Copy Data Management DellEMC Unity Restore job. Volumes can be restored to their original location or a new volume in the same or different DellEMC Unity storage system.

BEFORE YOU BEGIN:

- Create and run one or more DellEMC Unity Backup jobs. See [Create a Backup Job Definition - DellEMC Unity](#) on page **220**.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page **75**.

CONSIDERATIONS:

- Note that before running replication jobs, replication connections must be established between VNX arrays. Create replication connections through the DellEMC Unisphere wizard found under **Hosts > Replication Connections**.
- Note that to restore data to an original volume, you must first offline the target disk on the host prior to recovery. Once recovery completes, bring the target disk back online.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page **155**.

To create an Instant Disk Restore DellEMC Unity Restore job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **DellEMC Unity**.
2. Click **New** , then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Instant Disk Restore**  template.
5. Click **Source** . From the drop-down menu select **LUNs** or **File Systems**. Select a source site and an associated DellEMC Unity source to view volumes with available recovery points. Select resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **LUN Search** or **File System Search** from the drop-down menu to search for resources with available recovery points. Add copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.

Note: When selecting a specific version, data created through VMware Backup jobs that apply to the selected DellEMC Unity resource display, as the same data is contained with the snapshot for VMware and non-VMware related data.

7. Click **Destination** . Select the DellEMC Unity hosts that contain the iSCSI Qualified Name (IQN) or Fibre Channel WWPN of the application that you want to assign to.

Note: The DellEMC Unity hosts that are used during runtime may be different based on the initiator name.

8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enable to allow a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Session will auto cleanup/end after postscript execution completes

Enable to automatically clean up allocated resources after a post-script defined in the Post-Script field completes.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334.

State and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

- Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

- Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.
- When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

To create a Restore Volume(s) DellEMC Unity Restore job definition:

- Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **DellEMC Unity**.
- Click **New** , then select **Restore**. The job editor opens.
- Enter a name for your job definition and a meaningful description.
- Select the **Restore Volume(s)**  template.
- Click **Source** . From the drop-down menu select **LUNs** or **File Systems**. Select a source site and an associated DellEMC Unity source to view volumes with available recovery points. Select one or more

resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **LUN Search** or **File System Search** from the drop-down menu to search for resources with available recovery points. Add copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Destination** . To restore to the original volume, select **Restore to original volume**, or select **Restore to alternative location** and select a volume and associated pool. If no pool is selected, the pool with the largest amount of space available is chosen by default.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Overwrite volume if exists

Enable to overwrite the volume if the volume exists on the destination.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page **553**.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page **334**. **State** and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

- Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

- Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

- When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Jobs Overview](#) on page **161**
- [Using State and Status Arguments in Postscripts](#) on page **334**
- [Edit a Job Definition](#) on page **342**

- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**
- [Search and Filter Guidelines](#) on page **548**

Create a Restore Job Definition - IBM Spectrum Accelerate

IBM Spectrum Copy Data Management leverages Copy Data Management technology for recovering IBM Spectrum Accelerate volumes through IBM Spectrum Accelerate Restore jobs.

Instant Disk Restore

Provides instant writable access to a volume. An IBM Spectrum Copy Data Management snapshot is mapped to a target server where it can be accessed, copied, or put immediately into production use as needed.

Restore Volume(s)

Recover a volume from a FlashCopy or Global Mirror created through an IBM Spectrum Copy Data Management IBM Spectrum Accelerate Backup job. Volumes can be restored to their original location or a new volume in the same or different IBM storage system.

BEFORE YOU BEGIN:

- Create and run one or more IBM Spectrum Accelerate Backup jobs. See [Create a Backup Job Definition - IBM Spectrum Accelerate](#) on page 223.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- IBM providers utilize port 22 for communication with IBM Spectrum Copy Data Management.
- Note that to restore data to an original volume, you must first offline the target disk on the host prior to recovery. Once recovery completes, bring the target disk back online.
- Note that after restoring data to an alternate location you must map the host to the restore volume on the IBM storage system. Then rescan the disk on the host, and bring the disk online.
- In IBM storage environments, port grouping and IP partnerships are required to enable remote copy connections. See IBM's [SAN Volume Controller and Storwize Family Native IP Replication Guide](#).
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create an Instant Disk Restore IBM Spectrum Accelerate Restore job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **IBM Spectrum Accelerate**.
2. Click **New** , then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.

4. Select the **Instant Disk Restore**  template.
5. Click **Source** . From the drop-down menu select **Volume** to select a source site and an associated IBM source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.
Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.
6. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
Note: When selecting a specific version, data created through VMware Backup jobs that apply to the selected IBM resource display, as the same data is contained with the snapshot for VMware and non-VMware related data.
7. Click **Destination** . Select the IBM hosts that contain the iSCSI Qualified Name (IQN) or Fibre Channel WWPN of the application that you want to assign to.
Note: The IBM hosts that are used during runtime may be different based on the initiator name.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334. **State** and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.
Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.
12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

To create a Restore Volume(s) IBM Spectrum Accelerate Restore job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **IBM Spectrum Accelerate**.
2. Click **New** , then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Restore Volume(s)**  template.

5. Click **Source** . From the drop-down menu select **Volume** to select a source site and an associated IBM source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.
Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.
6. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Destination** . To restore to the original volume, select **Restore to original volume**, or select **Restore to alternative location** and select a volume and associated pool. If no pool is selected, the pool with the largest amount of space available is chosen by default.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334. **State** and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

- Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

- Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

- When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page 162.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page 164.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Create a Backup Job Definition - IBM Spectrum Accelerate](#) on page 223
- [Using State and Status Arguments in Postscripts](#) on page 334
- [Edit a Job Definition](#) on page 342
- [Delete a Job Definition](#) on page 343

- [Create a Schedule](#) on page **155**
- [Search and Filter Guidelines](#) on page **548**

Create a Restore Job Definition - IBM Spectrum Virtualize

IBM Spectrum Copy Data Management leverages Copy Data Management technology for recovering IBM Spectrum Virtualize volumes through IBM Spectrum Virtualize Restore jobs.

Instant Disk Restore

Provides instant writable access to a volume. An IBM Spectrum Copy Data Management snapshot is mapped to a target server where it can be accessed, copied, or put immediately into production use as needed.

Restore Volume(s)

Recover a volume from a FlashCopy or Global Mirror created through an IBM Spectrum Copy Data Management IBM Spectrum Virtualize Backup job. Volumes can be restored to their original location or a new volume in the same or different IBM storage system.

BEFORE YOU BEGIN:

- Create and run one or more IBM Spectrum Virtualize Backup jobs. See [Create a Backup Job Definition - IBM Spectrum Virtualize](#) on page **227**.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page **75**.

CONSIDERATIONS:

- IBM providers utilize port 22 for communication with IBM Spectrum Copy Data Management.
- Note that to restore data to an original volume, you must first offline the target disk on the host prior to recovery. Once recovery completes, bring the target disk back online.
- Note that after restoring data to an alternate location you must map the host to the restore volume on the IBM storage system. Then rescan the disk on the host, and bring the disk online.
- In IBM storage environments, port grouping and IP partnerships are required to enable remote copy connections. See IBM's [SAN Volume Controller and Storwize Family Native IP Replication Guide](#).
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page **155**.

To create an Instant Disk Restore IBM Spectrum Virtualize Restore job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **IBM Spectrum Virtualize**.
2. Click **New** , then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.

4. Select the **Instant Disk Restore**  template.
5. Click **Source** . From the drop-down menu select **Volume** to select a source site and an associated IBM source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.
6. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.

Note: When selecting a specific version, data created through VMware Backup jobs that apply to the selected IBM resource display, as the same data is contained with the snapshot for VMware and non-VMware related data.
7. Click **Destination** . Select the IBM hosts that contain the iSCSI Qualified Name (IQN) or Fibre Channel WWPN of the application that you want to assign to.

Note: The IBM hosts that are used during runtime may be different based on the initiator name.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Make IA clone resource permanent

Enable to turn the snapshot copy into a proper resource that will not be cleaned up after the Instant Disk Restore job completes.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334. **State** and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.
Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.
12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

To create a Restore Volume(s) IBM Spectrum Virtualize Restore job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **IBM Spectrum Virtualize**.
2. Click **New** , then select **Restore**. The job editor opens.

3. Enter a name for your job definition and a meaningful description.
4. Select the **Restore Volume(s)**  template.
5. Click **Source** . From the drop-down menu select **Volume** to select a source site and an associated IBM source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.
Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.
6. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Destination** . To restore to the original volume, select **Restore to original volume**, or select **Restore to alternative location** and select a volume and associated pool. If no pool is selected, the pool with the largest amount of space available is chosen by default.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page **553**.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page **334**. **State** and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Create a Backup Job Definition - IBM Spectrum Virtualize](#) on page **227**
- [Using State and Status Arguments in Postscripts](#) on page **334**
- [Edit a Job Definition](#) on page **342**

- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**
- [Search and Filter Guidelines](#) on page **548**

Create a Restore Job Definition - NetApp ONTAP

IBM Spectrum Copy Data Management leverages Copy Data Management technology for recovering NetApp ONTAP volumes and files through NetApp ONTAP Restore jobs.

Instant Disk Restore

Provides instant writable access to volume or LUN. An IBM Spectrum Copy Data Management snapshot is mapped to a target server where it can be accessed, copied, or put immediately into production use as needed.

Restore Volume(s)

Recover a volume from a primary snapshot, vault, or mirror copy created through an IBM Spectrum Copy Data Management NetApp ONTAP Backup job. Volumes can be restored to their original location or a new volume in the same or different NetApp ONTAP cluster or server.

Note that Restore Volume jobs are not available for NetApp ONTAP storage systems operating in 7-mode. 7-mode resources will not display in Source or Destination steps.

Restore File(s)

Recover files from a primary snapshot created through an IBM Spectrum Copy Data Management NetApp ONTAP Backup job. Files are restored to their original location.

BEFORE YOU BEGIN:

- Create and run one or more NetApp ONTAP Backup jobs. See [Create a Backup Job Definition - NetApp ONTAP](#) on page 231.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- Target volumes and datastores can be automatically expanded through the autogrow feature if supported by the underlying storage.
- Note that a volume restore through a NetApp ONTAP Restore job that includes NetApp ONTAP storage systems operating in 7-mode is not supported. These providers will not display during job creation.
- Note that NetApp ONTAP Backup jobs can only vault or mirror snapshots created through IBM Spectrum Copy Data Management jobs.
- Note that a file restore from a mirror location is not available for NetApp ONTAP storage systems operating in 7-mode.

- Restore Volume jobs are not available for NetApp ONTAP storage systems operating in 7-mode. 7-mode resources will not display in Source or Destination steps.
- Note that a file restore through a NetApp ONTAP Restore job can only utilize the alternate location feature if both the source and the destination are NetApp ONTAP storage systems running Command Data ONTAP 8.3.
- Note that NetApp ONTAP storage systems operating in 7-mode support file recovery from primary snapshots to their original locations. To restore files from a mirror source, create and run an Instant Disk Restore job with the mirror as a source, then mount the restored volume via CIFS or NFS. Files can then be copied to a new location.
- Note that the .snapshot folder must be visible on NFS shares in order to properly view and run Copy Data Management jobs on NetApp ONTAP storage systems running Data ONTAP in 7-Mode or Clustered Data ONTAP up to and including version 8.2. Confirm with your administrator that the .snapshot folder is not hidden in your NetApp ONTAP environment.
- NetApp ONTAP and VMware Restore jobs will fail if the iSCSI Initiator Group (iGroup) is not configured on the NetApp Clustered Data ONTAP 8.3 storage system target. The procedure only needs to be performed once. Previously created iGroups for earlier versions of NetApp Clustered Data ONTAP do not need to be reconfigured for version 8.3. Note that there should only be one iGroup using the software iSCSI initiator. For more information, contact Technical Support.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create an Instant Disk Restore NetApp ONTAP Restore job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **NetApp ONTAP** .
2. Click **New** , then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Instant Disk Restore**  template.
5. Click **Source** . From the drop-down menu select **Volume** to select a source site and an associated NetApp ONTAP source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid. Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.
6. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version**

field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.

Note: When selecting a specific version, data created through VMware Backup jobs that apply to the selected NetApp ONTAP resource display, as the same data is contained with the snapshot for VMware and non-VMware related data.

7. Click **Destination** . Select your NFS and CIFS mapping options, including the Volume Name Prefix. If a prefix is defined, the resulting NFS path or CIFS share displays as follows:

NFS: "/" + "volumeNamePrefix" + "_" + "sourcevolumeName"

CIFS: "volumeNamePrefix" + "_" + "sourcevolumeName"

If a prefix is not defined, a unique naming convention is applied. Unique prefixes should be defined for jobs that run concurrently. If a job is run with a specified prefix and the resulting Instant Disk Restore volume is made permanent, you must rename the NFS/CIFS path for the permanent volume prior to running the same job again.

8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Make IA clone resource permanent

Enable to turn the snapshot copy into a proper resource that will not be cleaned up after the instant access job completes.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the NetApp ONTAP volume recovery fails.

Allow to overwrite and force clean up of pending old session

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334. **State** and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.
Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.
12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

To create a Restore Volume(s) NetApp ONTAP Restore job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **NetApp ONTAP** .
2. Click **New** , then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Restore Volume(s)**  template.

5. Click **Source** . From the drop-down menu select **Volume** to select a source site and an associated NetApp ONTAP source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid. Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.
6. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Destination** . To restore to the original volume, select **Restore to original volume**, or select **Restore to new volume in the same or different NetApp ONTAP cluster or server** and select a volume and associated aggregate. If no aggregate is selected, the aggregate with the largest amount of space available is chosen by default.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Auto mount NFS after volume restored

Enable to automatically mount the restored volume after restoration completes.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or

create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334.

State and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

- Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

- Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.
- When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

To create a Restore File(s) NetApp ONTAP Restore job definition:

- Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **NetApp ONTAP** .
- Click **New** , then select **Restore**. The job editor opens.
- Enter a name for your job definition and a meaningful description.
- Select the **Restore File(s)**  template.
- Click **Source** . Select a source site and an associated NetApp ONTAP source to view volumes with available recovery points. Select recovery points and files to recover. Selected files are added to the **Selected Files** pane.
- Click **Copy** . Sites containing copies of the selected files display. Select a site. The latest copy of the file is used. If recovery from one snapshot fails, another copy from the same site is used.

7. Click **Destination** . To restore to the original volume, select **Restore to original volume**, or select **Restore to alternate location** and select an alternate volume or directory.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the file recovery fails.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334.

State and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.
Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.
12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Create a Backup Job Definition - NetApp ONTAP](#) on page **231**
- [Using State and Status Arguments in Postscripts](#) on page **334**
- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**
- [Search and Filter Guidelines](#) on page **548**

Create a Restore Job Definition - HPE Nimble Storage

IBM Spectrum Copy Data Management leverages Copy Data Management technology for recovering HPE Nimble Storage volumes through HPE Nimble Storage Restore jobs.

Instant Disk Restore

Provides instant writable access to a volume. An IBM Spectrum Copy Data Management snapshot is mapped to a target server where it can be accessed, copied, or put immediately into production use as needed.

Restore Volume(s)

Recover a volume from a snapshot created through an IBM Spectrum Copy Data Management HPE Nimble Storage Backup job. Volumes can be restored to their original location or a new volume in the same or different HPE Nimble Storage system.

BEFORE YOU BEGIN:

- Create and run one or more HPE Nimble Storage Backup job. See [Create a Backup Job Definition - HPE Nimble Storage](#) on page 234.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create an Instant Disk Restore HPE Nimble Storage Restore job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **HPE Nimble Storage**.
2. Click **New** , then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Instant Disk Restore**  template.
5. Click **Source** . From the drop-down menu select **Volume** to select a source site and an associated HPE Nimble Storage source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid. Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.
6. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Destination** . Select a HPE Nimble Storage destination.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334. **State** and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

- When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

To create a Restore Volume(s) HPE Nimble Storage Restore job definition:

- Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **HPE Nimble Storage**.
- Click **New** , then select **Restore**. The job editor opens.
- Enter a name for your job definition and a meaningful description.
- Select the **Restore Volume(s)**  template.
- Click **Source** . From the drop-down menu select **Volume** to select a source site and an associated HPE Nimble Storage source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid. Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.
- Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
- Click **Destination** . To restore to the original volume, select **Restore to original volume**, or select **Restore to alternative location** and select a volume and associated pool. If no pool is selected, the pool with the largest amount of space available is chosen by default.
- To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
- To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334. **State** and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page 162.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page 164.

- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Create a Backup Job Definition - HPE Nimble Storage](#) on page **234**
- [Using State and Status Arguments in Postscripts](#) on page **334**
- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**
- [Search and Filter Guidelines](#) on page **548**

Create a Restore Job Definition - Pure Storage FlashArray

IBM Spectrum Copy Data Management leverages Copy Data Management technology for recovering Pure Storage FlashArray volumes through Pure Storage FlashArray Restore jobs.

Instant Disk Restore

Provides instant writable access to a volume. An IBM Spectrum Copy Data Management snapshot is mapped to a target server where it can be accessed, copied, or put immediately into production use as needed.

Restore Volume(s)

Recover a volume from a snapshot or replication created through an IBM Spectrum Copy Data Management Pure Storage FlashArray Backup job. Volumes can be restored to their original location or a new volume in the same or different Pure Storage system.

Pure Storage FlashArray can be set to utilize CloudSnap functionality so that snapshots are offloaded from the local storage array to cloud storage. When a recovery job runs, the snapshot will be recovered from the local storage array. If it has been condensed out of the local storage and the **Add Snapshot Offload**  option was set during SLA creation, the offload copy from cloud storage will be restored.

Best Practice: The ECX user interface will indicate that a Pure Storage FlashArray CloudSnap offload job has completed even though the transfer is still occurring in the background which is dependent on network speeds. Consider setting an age as the retention for offload copies when using the Pure Storage FlashArray CloudSnap functionality. Doing so will ensure that a sufficient amount of time has passed for data to be transferred to the S3 storage target before it has to be condensed out from a backup. This is particularly important if several offload jobs are run in quick succession.

BEFORE YOU BEGIN:

- Create and run one or more Pure Storage FlashArray Backup job. See [Create a Backup Job Definition - Pure Storage FlashArray](#) on page 238.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.

CONSIDERATIONS:

- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

To create an Instant Disk Restore Pure Storage FlashArray Restore job definition:

1. Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **Pure Storage FlashArray**.
2. Click **New** , then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Instant Disk Restore**  template.

5. Click **Source** . From the drop-down menu select **Volume** to select a source site and an associated Pure Storage source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid. Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.
6. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Destination** . Select a Pure Storage FlashArray destination.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334. **State** and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

- Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

- Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.
- When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

To create a Restore Volume(s) Pure Storage FlashArray Restore job definition:

- Click the **Jobs**  tab. Expand the **Storage Controller** folder, then select **Pure Storage FlashArray**.
- Click **New** , then select **Restore**. The job editor opens.
- Enter a name for your job definition and a meaningful description.
- Select the **Restore Volume(s)**  template.
- Click **Source** . From the drop-down menu select **Volume** to select a source site and an associated Pure Storage source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.
- Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version**

field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.

7. Click **Destination** . To restore to the original volume, select **Restore to original volume**, or select **Restore to alternative location** and select a volume and associated pool. If no pool is selected, the pool with the largest amount of space available is chosen by default.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334.

State and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Create a Backup Job Definition - Pure Storage FlashArray](#) on page **238**
- [Using State and Status Arguments in Postscripts](#) on page **334**
- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**
- [Search and Filter Guidelines](#) on page **548**

Create a Restore Job Definition - VMware

IBM Spectrum Copy Data Management leverages Copy Data Management technology for testing and cloning use cases, instant recovery, and full disaster recovery. VMware Backup jobs support Instant VM Restore and Instant Disk Restore scenarios.

Instant VM Restore  jobs are run in the following modes:

Test Mode

Creates temporary virtual machines for development/testing, snapshot verification, and disaster recovery verification on a scheduled, repeatable basis without affecting production environments. Test machines are kept running as long as needed to complete testing and verification and are then cleaned up after testing and verification completes. Through fenced networking, you can establish a safe environment to test your jobs without interfering with virtual machines used for production. Virtual machines created through Test mode are also given unique names and identifiers to avoid conflicts within your production environment.

Clone Mode

Creates copies of virtual machines for use cases requiring permanent or long-running copies for data mining or duplication of a test environment in a fenced network. Virtual machines created through Clone mode are also given unique names and identifiers to avoid conflicts within your production environment. With clone mode you must be sensitive to resource consumption, since clone mode creates permanent or long-term virtual machines.

Production Mode

Enables disaster recovery at the local site from primary storage or a remote disaster recovery site, replacing original machine images with recover images. All configurations are carried over as part of the recovery, including names and identifiers, and all copy data jobs associated with the virtual machine continue to run.

You can also set an IP address or subnet mask for virtual machines to be repurposed for development/testing or disaster recovery use cases. Supported mapping types include IP to IP, IP to DHCP, and subnet to subnet.

Instant Disk Restore

Provides instant writable access to data and application recovery points. An IBM Spectrum Copy Data Management snapshot is mapped to a target server where it can be accessed, copied, or put immediately into production use as needed.

BEFORE YOU BEGIN:

- Create and run a VMware Backup job. See [Create a Backup Job Definition - VMware](#) on page 242.

- Ensure the latest version of VMware Tools is installed in your environment. IBM Spectrum Copy Data Management was tested against VMware Tools 9.10.0.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page 75.
- You must add credentials to the destination virtual machine when recovering with the subnet option. See [Add Credentials to a Virtual Machine](#) on page 101.
- To run network prescripts and postscripts, select the **Power on after recovery option**, then the **Enable pre and post network VM level scripts option**. Note that prior to running the job, the scripts must be defined and copied to specific locations on the virtual machine. In a Windows environment, copy prenetwork.bat and postnetwork.bat to your c:\program files\ibm\IBM Spectrum Copy Data Management\scripts directory. In a Linux environment, copy prenetwork.sh and postnetwork.sh to your /opt/CDM/scripts/ directory.

CONSIDERATIONS:

- Note that VMware Backup and Restore jobs only support vCenters or ESX hosts running vSphere 5.1 through 6.5.
- If a recovery through an Instant VM Restore RRP job using an AWS-based SLA Policy fails with a *VMwareVMotionException, A general system error occurred: The source detected that the destination failed to resume.* error, retry the job at a later time. This error may display if transferring data from the AWS cloud takes too long.
- When running VADP-based VM Replication workflows, target volumes and datastores can be automatically expanded in response to space usage requirements if supported by the underlying storage. Automatic growing prevents a volume from running out of space or forcing you to delete files manually. For a list of supported storage systems see [System Requirements](#) on page 23.
- Note that VMware DRS cluster datastores are supported in VMware Backup and Restore jobs.
- Note that after an Instant Disk Restore Restore job completes, your vDisk will be mounted but you may need to bring it online through the operating system from the Disk Management console.
- In addition to NFS, IBM Spectrum Copy Data Management supports VMFS datastores for NetApp ONTAP storage targets.
- Note that Instant Disk Restore recoveries utilizing the VM Replication method are not supported at the datastore level. Instant Disk Restore datastore level recoveries are supported through the primary storage snapshot method.

- Note that in Instant VM Restore recoveries utilizing NetApp ONTAP storage systems running Clustered Data ONTAP or Data ONTAP operating in 7-Mode, if a source with a swap directory on a dedicated datastore is recovered to a different destination, then the source datastore must have more free space than the amount of memory configured for the virtual machine. This may not be applicable if the virtual machine is configured with memory reservation.
- Instant Disk Restore recoveries of VMDKs through snapshots of IBM Spectrum Protect Snapshot-protected virtual machines are not supported.
- Instant VM Restore recoveries of IBM Spectrum Protect Snapshot-protected virtual machines under vApps are only restored as standalone virtual machines and not under or in association with the vApp.
- A VMware Restore job recovering a virtual machine from an ESX cluster protected with snapshot, vault, or mirror displays a Locate LUN failure if the maximum allowed LUNs for the ESX host recovery target reaches its limit.
- In NetApp ONTAP environments running Clustered Data ONTAP, cluster peering must be enabled. Peer relationships enable communication between SVMs. See NetApp ONTAP's [Cluster and Vserver Peering Express Guide](#).
- In IBM storage environments, port grouping and IP partnerships are required to enable remote copy connections. See IBM's [SAN Volume Controller and Storwize Family Native IP Replication Guide](#).
- NetApp ONTAP and VMware Restore jobs will fail if the iSCSI Initiator Group (iGroup) is not configured on the NetApp Clustered Data ONTAP 8.3 storage system target. The procedure only needs to be performed once. Previously created iGroups for earlier versions of NetApp Clustered Data ONTAP do not need to be reconfigured for version 8.3. For more information, contact Technical Support.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 155.

CONSIDERATIONS FOR VMWARE VIRTUAL VOLUMES:

- All functionality of Restore workflows are supported where the original virtual machine was stored on a virtual volume (VVOL) datastore. Virtual machines can be recovered to a VVOL target or VMware datastore through Production or Clone mode.
- With the introduction of VVOLs, a storage vendor's vSphere API for Storage Awareness (VASA) stores metadata required to run virtual machines on VVOL datastores. IBM Spectrum Copy Data Management does not require the VASA provider metadata for Restore workflows, as IBM Spectrum Copy Data Management stores the VM Replication in a VMware target datastore. In case of a disaster in which the VASA provider is lost, a

new VASA provider can be brought up to recover the virtual machine back to a VVOL datastore through Production or Clone mode.

Best Practice: Create a schedule before creating a job definition so that you can easily add the schedule to the job definition.

To create an Instant Disk Restore VMware Restore job definition:

1. Click the **Jobs**  tab. Expand the **Hypervisor** folder, then select **VMware**.
2. Click **New** , then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Instant Disk Restore**  template.
5. Click **Source** , then select **VM Storage** or **Datastores** as the source type. Select a source site and an associated VMware source to view virtual machines, VM templates, folders, vApps, and datacenters with available recovery points. Select resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.
6. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Destination** . Expand a VMware source to view virtual machines, folders, vApps, and datacenters available as destinations. To restore to the original host or cluster, select **Use original host or cluster**.
8. Select the datastore and virtual disk mapping options if you selected a destination different from the original host or cluster.

Virtual Disks

In the **VM** field, select virtual machine destinations.

In the **Disk Mode** field, select “persistent”, “independent persistent”, or “independent non-persistent”.

The values are defined as follows:

- **Persistent** - Changes are permanently written to the virtual disk. The disk is included in snapshots taken of its virtual machine.
- **Independent Persistent** - Changes are permanently written to the virtual disk. The disk is excluded from any snapshots taken of its virtual machine.
- **Independent Non-Persistent** - Changes to the virtual disk are discarded when the virtual machine powers off; the VMDK files revert to their original state. The disk is excluded from any snapshots taken of its virtual machine.

In the optional **Controller Type** field, select a supported SCSI controller, including LSI SAS, LSI Parallel, BusLogic, and VMware Paravirtual. Changing the SCSI controller type replaces the existing controller

with a new controller, applies the common settings of the existing controller to the new controller, and reassigns all SCSI devices to the new controller.

Use the optional **Controller Address #** and **Controller LUN #** fields to select specific controllers or LUNs.

Datstores

Set the destination datastore.

9. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs**  tab.
10. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Protocol Priority

If more than one storage protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

Make IA clone resource permanent

Enable to turn the snapshot copy into a proper resource that will not be cleaned up after the instant access job completes.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the virtual machine recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page **152**.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or

create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334. **State** and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

11. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

12. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

13. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

To create an Instant VM Restore VMware Restore job definition:

1. Click the **Jobs**  tab. Expand the **Hypervisor** folder, then select **VMware**.
2. Click **New** , then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Instant VM Restore**  template.
5. Click **Source** . From the drop-down menu, select **VMs and Templates** to choose a source site and an associated VMware source to view virtual machines, VM templates, datastores, folders, and vApps and with available recovery points. Select resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

- Alternatively, select **VM Search** from the drop-down menu to search for virtual machines with available recovery points across all datacenters. Add virtual machine copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.
6. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
 7. Click **Destination** . Select a destination site and an associated VMware destination to view virtual machines, folders, vApps, and datacenters available as destinations. To restore to the original host or cluster and allow IBM Spectrum Copy Data Management to define the destination IP address, select **Use original host or cluster with system defined IP configuration**. To restore to the original host or cluster using your predefined IP address configuration, select **Use original host or cluster with original IP configuration**. To restore to a destination different from the original host or cluster, select **Use alternative host or cluster**.
 8. Select virtual network and datastore mapping options if you selected **Use alternative original host or cluster** in the previous step. The Virtual Networks pane displays all of the virtual networks associated with your VMware Restore job sources. New virtual networks must be selected for use at the recovery site, as well as new datastores on the Datastores pane. Select a production and test network in the Virtual Networks tab, and a destination datastore in the Datastore tab.

Virtual Networks

Set virtual networks for production and test recovery jobs. Destination network settings for production and test environments should be different locations.

Note: Network mappings are disabled for IBM Spectrum Protect Snapshot-protected virtual machines in Instant VM Restore restores to alternate locations. You must re-enable virtual machine networks to the proper target networks manually through vCenter once restoration completes.

Datastores

Set the destination datastore.

Subnet

Set an IP address or subnet mask for virtual machines to be repurposed for development/testing or disaster recovery use cases. Supported mapping types include IP to IP, IP to DHCP, and subnet to subnet. Virtual machines containing multiple NICs are supported.

By default, the **Use system defined subnets and IP addresses for VM guest OS on destination** option is enabled. To use your predefined subnets and IP addresses, select **Use original subnets and IP addresses for VM guest OS on destination**.

To create a new mapping configuration, select **Add mappings for subnets and IP addresses for VM guest OS on destination**, then click **Add Mapping**. Enter a subnet or IP address in the Source field. In

the destination field, select **DHCP** to automatically select an IP and related configuration information if DHCP is available on the selected client. Select **Static** to enter a specific subnet or IP address, subnet mask, gateway, and DNS. Note that **Subnet or IP Address**, **Subnet Mask**, and **Gateway** are required fields. If a subnet is entered as a source, a subnet must also be entered as a destination.

IP reconfiguration is skipped for virtual machines if a static IP is used but no suitable subnet mapping is found, or if the source machine is powered off and there is more than one associated NIC. In a Windows environment, if a virtual machine is DHCP only, then IP reconfiguration is skipped for that virtual machine. In a Linux environment all addresses are assumed to be static, and only IP mapping will be available.

Note: You must add credentials to the destination virtual machine when recovering with the subnet option. Note that if using a domain user account, the credentials must be added to the destination virtual machine, then configured through the **Test & Configure** option. This option verifies communication with the server, tests DNS settings between the IBM Spectrum Copy Data Management appliance and the server, and installs an IBM Spectrum Copy Data Management agent on the server. From the **Configure**  tab in the Provider Browser pane, right-click the virtual machine, then click **Test & Configure** . See [Add Credentials to a Virtual Machine](#) on page 101 and [Register a Provider](#) on page 75.

VM Folder

Set the VM folder path on the destination datacenter. To use an alternate path, deselect **Use original VM folder path on destination** and enter the path in the **Use alternative path** field. Note that the directory will be created if it does not exist. Use "/" as the root VM folder of the targeted datacenter.

9. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
10. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Default Mode

Set the VMware Restore job to run in Test, Production, or Clone mode by default. Once the job is created, it can be run in Test, Production, or Clone mode through the **Jobs**  tab.

Protocol Priority

If more than one storage protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

Power on after recovery

Toggle the power state of a virtual machine after a recovery is performed. Virtual machines are powered on in the order they are recovered, as set in the Source step. Turning this feature on also gives access to the **Enable pre and post network VM level scripts** option. Note that restored VM templates cannot be powered on after recovery.

Enable pre and post network VM level scripts

If the **Power on after recovery option** is enabled, network prescripts and postscripts can be run at the virtual machine level. Note that prior to running a VMware Backup job, the scripts must be defined and

copied to specific locations on the virtual machine. In a Windows environment, copy `prenetwork.bat` and `postnetwork.bat` to your `c:\program files\ibm\IBM Spectrum Copy Data Management\scripts` directory. In a Linux environment, copy `prenetwork.sh` and `postnetwork.sh` to your `/opt/CDM/scripts/` directory.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the virtual machine recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334.

State and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

11. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

12. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.
13. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.
14. Once the job completes successfully, select one of the following options from the **Actions** menu on the General tab of the job session on the Jobs tab: **End IV (Cleanup)**, **RRP (vMotion)**, or **Clone (vMotion)**.

End IV (Cleanup) destroys the virtual machine and cleans up all associated resources. Since this is a temporary/testing virtual machine, all data is lost when the virtual machine is destroyed.

RRP (vMotion) is equivalent to using the Production selection in the job Advanced screen. This option migrates the virtual machine through vMotion to the Datastore and the Virtual Network defined as the "For Production" Network.

Clone (vMotion) is equivalent to using the Clone selection in the job Advanced screen. This option migrates the virtual machine through vMotion to the Datastore and Virtual Network defined as the "For Test" network.

To create a Instant VM Restore (Long Distance) VMware Restore job definition:

1. Click the **Jobs**  tab. Expand the **Hypervisor** folder, then select **VMware**.
2. Click **New** , then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Instant VM Restore (Long Distance)**  template.
5. Click **Source** . From the drop-down menu, select **VMs and Templates** to choose a source site and an associated VMware source to view virtual machines, VM templates, datastores, folders, and vApps and with available recovery points. Select resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **VM Search** from the drop-down menu to search for virtual machines with available recovery points across all datacenters. Add virtual machine copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy** . Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Local Destination** . Select a local destination site and an associated VMware destination to view virtual machines, folders, vApps, and datacenters available as destinations. Note that the local storage destination is temporary, and is cleaned up after the job completes.
8. Click **Remote Destination** . Select a remote destination site and an associated VMware destination to view virtual machines, folders, vApps, and datacenters available as destinations.
9. Select virtual network and datastore mapping options. The Virtual Networks pane displays all of the virtual networks associated with your VMware Restore job sources. New virtual networks must be selected for use at the recovery site, as well as new datastores on the Datastores pane. Select a production and test network in the Virtual Networks tab, and a destination datastore in the Datastore tab.

Virtual Networks

Set virtual networks for production and test recovery jobs. Destination network settings for production and test environments should be different locations.

Note: Network mappings are disabled for IBM Spectrum Protect Snapshot-protected virtual machines in Instant VM Restore restores to alternate locations. You must re-enable virtual machine networks to the proper target networks manually through vCenter once restoration completes.

Datastores

Set the destination datastore.

Subnet

Set an IP address or subnet mask for virtual machines to be repurposed for development/testing or disaster recovery use cases. Supported mapping types include IP to IP, IP to DHCP, and subnet to subnet. Virtual machines containing multiple NICs are supported.

By default, the **Use system defined subnets and IP addresses for VM guest OS on destination** option is enabled. To use your predefined subnets and IP addresses, select **Use original subnets and IP addresses for VM guest OS on destination**.

To create a new mapping configuration, select **Add mappings for subnets and IP addresses for VM guest OS on destination**, then click **Add Mapping**. Enter a subnet or IP address in the Source field. In the destination field, select **DHCP** to automatically select an IP and related configuration information if DHCP is available on the selected client. Select **Static** to enter a specific subnet or IP address, subnet mask, gateway, and DNS. Note that **Subnet or IP Address**, **Subnet Mask**, and **Gateway** are required fields. If a subnet is entered as a source, a subnet must also be entered as a destination.

IP reconfiguration is skipped for virtual machines if a static IP is used but no suitable subnet mapping is found, or if the source machine is powered off and there is more than one associated NIC. In a Windows

environment, if a virtual machine is DHCP only, then IP reconfiguration is skipped for that virtual machine. In a Linux environment all addresses are assumed to be static, and only IP mapping will be available.

Note: You must add credentials to the destination virtual machine when recovering with the subnet option. Note that if using a domain user account, the credentials must be added to the destination virtual machine, then configured through the **Test & Configure** option. This option verifies communication with the server, tests DNS settings between the IBM Spectrum Copy Data Management appliance and the server, and installs an IBM Spectrum Copy Data Management agent on the server. From the **Configure**  tab in the Provider Browser pane, right-click the virtual machine, then click **Test & Configure** . See [Add Credentials to a Virtual Machine](#) on page 101 and [Register a Provider](#) on page 75.

VM Folder

Set the VM folder path on the destination datacenter. To use an alternate path, deselect **Use original VM folder path on destination** and enter the path in the **Use alternative path** field. Note that the directory will be created if it does not exist. Use "/" as the root VM folder of the targeted datacenter.

10. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
11. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Default Mode

Set the VMware Restore job to run in Production or Clone mode by default. Once the job is created, it can be run in Production or Clone mode through the **Jobs**  tab.

Protocol Priority

If more than one storage protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

Power on after recovery

Toggle the power state of a virtual machine after a recovery is performed. Virtual machines are powered on in the order they are recovered, as set in the Source step. Turning this feature on also gives access to the **Enable pre and post network VM level scripts** option. Note that restored VM templates cannot be powered on after recovery.

Enable pre and post network VM level scripts

If the **Power on after recovery option** is enabled, network prescripts and postscripts can be run at the virtual machine level. Note that prior to running a VMware Backup job, the scripts must be defined and copied to specific locations on the virtual machine. In a Windows environment, copy prenetwork.bat and postnetwork.bat to your c:\program files\ibm\IBM Spectrum Copy Data Management\scripts directory. In a Linux environment, copy prenetwork.sh and postnetwork.sh to your /opt/CDM/scripts/ directory.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the virtual machine recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts**  view on the **Configure**  tab. See [Configure Scripts](#) on page 152.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#) on page 553.

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#) on page 334.

State and **status** arguments are not supported for Backup jobs.

Select **Continue operation on script failure** to continue running the job if a command in any of the scripts associated with the job fails.

- Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

- Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view

the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

14. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.
15. Once the job completes successfully, select one of the following options from the **Actions** menu on the General tab of the job session on the Jobs tab: **End IV (Cleanup)**, **RRP (vMotion)**, or **Clone (vMotion)**.
RRP (vMotion) is equivalent to using the Production selection in the job Advanced screen. This option migrates the virtual machine through vMotion to the Datastore and the Virtual Network defined as the "For Production" Network.
Clone (vMotion) is equivalent to using the Clone selection in the job Advanced screen. This option migrates the virtual machine through vMotion to the Datastore and Virtual Network defined as the "For Test" network.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Create a Backup Job Definition - VMware](#) on page **242**
- [Using State and Status Arguments in Postscripts](#) on page **334**
- [NetApp ONTAP Document: Cluster and Vserver Peering Express Guide](#)
- [NetApp ONTAP Document: iSCSI Configuration and Provisioning for ESX Express Guide](#)
- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**
- [Search and Filter Guidelines](#) on page **548**

Restore Jobs - Rename Mount Points and Initialization Parameter Options

This option is available for supported Restore jobs running Instant Database Restore and DevOps workflows.

Mount Point Rename

This option controls how mount points and ASM diskgroups are renamed when they are mounted on the destination host/cluster during recovery.

Append a timestamp: By default, IBM Spectrum Copy Data Management appends a timestamp to the original mount point. For example:

Mount Point Rename - Append a timestamp

Original Mount Point	New Mount Point
/u02/myproddb/data	/u02/myproddb/data_1479400505
+MYPRODDATA	+MYPRODDATA1479400505

Do not rename: Select this option if you do not want to rename mount points or ASM diskgroups during recovery. IBM Spectrum Copy Data Management will mount them with the same path/name as the source.

Add a custom prefix: Select this option and specify a custom prefix to be prepended to the source paths/names. The prefix value may contain leading or trailing slashes. In the case of ASM diskgroup names, the slashes are removed. For example:

Mount Point Rename - Add a custom prefix

Original Mount Point	Prefix	New Mount Point
/u02/myproddb/data	tmp	/tmpu02/myproddb/data
/u02/myproddb/data	/tmp/	/tmp/u02/myproddb/data
+MYPRODDATA	/tmp/	+TMPMYPRODDATA

Add a custom suffix: Select this option and specify a custom suffix to be appended the source paths/names. For example:

Mount Point Rename - Add a custom suffix

Original Mount Point	Suffix	New Mount Point
/u02/myproddb/data	tmp	/u02/myproddb/datatmp

Original Mount Point	Suffix	New Mount Point
/u02/myproddb/data	/tmp/	/u02/myproddb/data/tmp
+MYPRODDATA	/tmp/	+MYPRODDATATMP

Replace a substring: Select this option and specify a custom string of characters in the old mount point to be replaced with another string of characters. The old and new substrings must be simple alphanumeric strings with no special characters. All occurrences of the specified substring are replaced. The value of the new substring may be left empty if you want to remove the old substring. For example:

Mount Point Rename - Replace a substring

Original Mount Point	Old Substring	New Substring	New Mount Point
/u02/myproddata/data	prod	test	/u02/mytestdata/data
/u02/myproddata/data	data	newdata	/u02/myprodnewdata/newdata
+MYPRODDATA	prod	test	+MYTESTDATA
+MYPRODDATA	prod	<empty>	+MYDATA

Database Initialization Parameters

This option controls the initialization parameters used to start up the recovered database in Oracle Instant Database Restore and DevOps workflows.

Use same parameters as source: This is the default option. IBM Spectrum Copy Data Management uses the same initialization parameters as the source database, but with the following changes:

- Parameters that contain paths such as *control_files*, *db_recovery_file_dest*, or *log_archive_dest_** are updated to reflect the new paths based on the renamed mount points of the recovered volumes.
- Parameters such as *audit_file_dest* and *diagnostic_dest* are updated to point to the appropriate location under the Oracle Base directory on the destination server if the path differs from the source server.
- The *db_name* and *db_unique_name* are updated to reflect the new name of the database if a new name is specified.
- Cluster-related parameters such as *instance_number*, *thread*, and *cluster_database* are set automatically by IBM Spectrum Copy Data Management depending on the appropriate values for the destination.

Use a template pfile: You can customize the initialization parameters by specifying a template file containing the initialization parameters that IBM Spectrum Copy Data Management should use.

The specified path must be to a plain text file that exists on the destination server and is readable by the IBM Spectrum Copy Data Management user. The file must be in Oracle pfile format, consisting of lines in the form `name = value`. Comments beginning with the `#` character are ignored.

IBM Spectrum Copy Data Management reads the template pfile and copies the entries to the new pfile that will be used to start up the recovered database. However, the following parameters in the template are ignored. Instead, IBM Spectrum Copy Data Management sets their values to reflect appropriate values from the source database or to reflect new paths based on the renamed mount points of the recovered volumes.

- `control_files`
- `db_block_size`
- `db_create_file_dest`
- `db_recovery_file_dest`
- `log_archive_dest`
- `spfile`
- `undo_tablespace`

Additionally, cluster-related parameters like *instance_number*, *thread*, and *cluster_database* are set automatically by IBM Spectrum Copy Data Management depending on the appropriate values for the destination.

RELATED TOPICS:

- [Create a Restore Job Definition - Oracle](#) on page **264**
- [Create a Restore Job Definition - InterSystems Caché](#) on page **252**
- [Oracle Requirements](#) on page **41**

Using State and Status Arguments in Postscripts

The functionality described in this topic (**state** and **status** arguments) applies only to postscripts for Restore jobs. Backup jobs are not supported. Prescripts are not supported.

Certain positional arguments can be passed to a Restore job postscript for conditional logic. The arguments that can be passed are **state** and **status**. You can use the **status** argument, for example, if you want to perform an action only if a Restore job completes successfully. Your script would perform the action only if the value of **status** were SUCCESS.

When passing the arguments via the **Postscript** field in the IBM Spectrum Copy Data Management user interface, the arguments must be surrounded by underscores (`_STATE_` and `_STATUS_`). When doing so, the arguments are replaced in the script with the actual corresponding values.

Specifically, to pass the arguments **state** and **status**, enter the following in the **Parameters** field of the Postscript section of a job:

```
_STATE_ _STATUS_
```

So for the above postscript example, the following output would display if the job ran in an Instant Disk Restore state and completed successfully:

```
state IA
```

```
status SUCCESS
```

Note: Passing `_STATE_` or `_STATUS_` to a script for an unsupported script type (such as prescripts) or an unsupported job type (such as Backup job scripts), simply passes “`_STATE_`” or “`_STATUS_`” as plain text.

An additional argument, `_VOLUMES_`, can also be added to a postscript to list the names of restored volumes.

RELATED TOPICS:

- [Create a Restore Job Definition - DellEMC Unity](#) on page **281**
- [Create a Restore Job Definition - IBM Spectrum Virtualize](#) on page **293**
- [Create a Restore Job Definition - NetApp ONTAP](#) on page **299**
- [Create a Restore Job Definition - VMware](#) on page **317**

System Jobs

The topics in the following section cover System job definitions as well as Maintenance job information.

Maintenance Job

The Maintenance job removes resources and associated objects created by IBM Spectrum Copy Data Management when a job in a pending state is deleted. The cleanup procedure reclaims space on your storage devices, cleans up your IBM Spectrum Copy Data Management catalog, and removes related snapshots. By default, the Maintenance job runs once a day, but the job's associated schedule can be altered to run more or less frequently depending on your needs, or the job can be run manually. The job cannot be deleted.

The Maintenance job only performs cleanup operations once a job in a pending state is deleted. All logs associated with the deleted job are removed from IBM Spectrum Copy Data Management, so it is advised to download job logs before the Maintenance job's next run. The job can be stopped and resumed; all pending operations set to occur before the job was stopped will resume upon the next job run.

After deleting a pending Application, DellEMC Unity, IBM, NetApp ONTAP, Pure Storage FlashArray, or VMware Backup job, all associated copy data, including recovery points, are deleted. The Maintenance job removes all VM Copies and Primary copies associated with deleted VMware Backup and Restore jobs. Similarly, after deleting a pending DellEMC Unity, IBM, NetApp ONTAP, or Pure Storage FlashArray Backup or Restore job, all associated DellEMC Unity, IBM, NetApp ONTAP, and Pure Storage FlashArray locations are removed by the Maintenance job. Once the Maintenance job completes, application, DellEMC Unity, IBM, NetApp ONTAP, Pure Storage FlashArray or VMware data that was copied as part of the backup job cannot be recovered. Any data related to the deleted job will not be recoverable.

The Maintenance job also removes cataloged data associated with deleted Application, DellEMC Unity, IBM, NetApp ONTAP, Pure Storage FlashArray and VMware Inventory jobs, and removes jobs and job sessions related to Script and Report jobs from the IBM Spectrum Copy Data Management interface.

RELATED TOPICS:

- [Delete a Job Definition](#) on page 343

Create a Report Job Definition

A Report job is a user-defined set of tasks and rules which run predefined or customized reports through a schedule that you define. The reports summarize information about cataloged providers and the data and other resources that reside on them. Reports generated during the job can be emailed in a variety of formats.

BEFORE YOU BEGIN:

- If you have a specific set of data that you want to report on, create a report with customized parameters to include in the Report job. See [Create a Customized Report](#) on page **363**.
- At least one provider must be associated with an Inventory job. Before defining an Inventory job, add providers. See [Register a Provider](#) on page **75**.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#) on page **75**.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page **155**.

To create a Report job definition:

1. Click the **Jobs**  tab. Expand the **System** folder, then select **Reports** .
2. Click **New** , then select **Analyze**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available reports, select one or more reports to include in the job definition. Expand reports to view associated customized reports. Select **Default** to run the predefined report parameters.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.
7. Optionally, select render options for your emailed report attachments. Reports can be rendered as Adobe PDFs, Microsoft Word files, and Microsoft Excel spreadsheets. Select **Export Data** to save reports to a defined location.
8. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

- Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.
Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.
- When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

Note: If you selected the **Start job now** option, the job runs.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Report Overview](#) on page **360**
- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**

Create a Script Job Definition

A Script job defines a set of commands to run on the IBM Spectrum Copy Data Management appliance. Use the Script job feature to add functionality to IBM Spectrum Copy Data Management.

A Script job includes a user-defined script as well as other job parameters.

A script can consist of one or many commands, such as a shell script. Command entries may include any entry that is valid for a Linux CentOS command shell. The author of the script must be familiar with the command, options, and parameters.

As the script runs, IBM Spectrum Copy Data Management interprets a return code of 0 as success and subsequent values as follows:

- Return codes 1-125 indicate a failure, the meaning of which is specific to the command or shell script invoked
- Return code 126 indicates the script command was found but is not executable
- Return code 127 indicates that the script command was not found
- Return codes 128-255 indicate that the command abnormally terminated. The formula 128+N is used with N representing the signal the process terminated on
- Values outside the range of 0-255, in both negative and positive values, are wrapped around the same series of return codes and are matched to a corresponding code. When positive error codes display over 255, they read as follows: 256=0, 257=1, 258=2, etc. When negative return codes display, the corresponding codes are as follows: -1 = 255, -2 = 254, -3 = 253, etc.

Example of a simple script:

Following is a sample one line script that runs a single command:

```
echo hello
```

Once this script is saved and subsequently run as a job, the word “hello” is output on the virtual appliance.

BEFORE YOU BEGIN:

- Ensure that the applications and programs called from your script are properly configured.
- Ensure that paths are properly set up.
- Ensure that authentication is properly set up.
- Ensure that schedules, SMTP providers, and other relevant providers are added in IBM Spectrum Copy Data Management. See [Create a Schedule](#) on page 155 and [Register a Provider](#) on page 75.

To create a Script job definition:

1. Click the **Jobs**  tab. Expand the **System** folder, then select **Scripts** .
2. Click **New** , then select **Script**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Enter the first command of your script in the command text box, and optionally enter a tag or remark in the corresponding description text box. Then click **Add** .
5. Continue to add commands. Use the reorder, edit, and delete options to assist you. Your script can consist of one or many commands.
6. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs**  tab.
7. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Run scripts in order

Select **Run scripts in order** for the commands in the job to run sequentially. The job runs the first command and when it finishes it runs the next one. Clear this option for the commands to run concurrently.

Stop execution on failure

Select **Stop execution on failure** for the job to stop running as soon as one of the commands fails. Clear this option for the job to continue to run after one of the commands fails. This option is only applicable if **Run scripts in order** is selected.

8. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add**  to add it to the list.

9. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.
Note: To create and select a new schedule, click the **Configure**  tab, then select **Schedules** . Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.
10. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs**  tab.

Note: If you selected the **Start job now** option, the job runs.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Return Code Reference](#) on page **553**
- [Edit a Job Definition](#) on page **342**
- [Delete a Job Definition](#) on page **343**
- [Create a Schedule](#) on page **155**

Edit a Job Definition

Revise a job definition to change the provider that the job is run against, SLA Policy associated with a job definition, options, notification and schedule properties.

BEFORE YOU BEGIN:

- Review the properties of your current job definitions. See [Jobs Overview](#) on page **161**.

To edit the properties of a job definition:

1. Click the **Jobs**  tab.
2. Select the job definition to edit by clicking in the row containing the job definition name.
3. Click **Edit** . The Job Definition Editor opens.
4. Click through the wizard, making revisions as needed.
5. Click **Finish** when you are satisfied that the job-specific information is correct.

NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#) on page **162**.
- Track the progress of the job on the Jobs tab. See [Monitor a Job Session](#) on page **164**.
- If SMTP options are enabled, an email message with information about the status of each task is sent when the job completes.

RELATED TOPICS:

- [Delete a Job Definition](#) on page **343**

Delete a Job Definition

Delete a job definition when it becomes obsolete. This keeps your operations current.

Note: The Maintenance job removes providers and associated objects created by IBM Spectrum Copy Data Management when a **Backup** or **Restore** job in a pending state is deleted. This job is run automatically.

To delete a job definition:

1. Click the **Jobs**  tab.
2. Select the job definition to delete by clicking in the row containing the job definition name.
3. Click **Delete** . A confirmation dialog box opens.
4. Confirm deletion. The job definition is deleted.

RELATED TOPICS:

- [Edit a Job Definition](#) on page **342**
- [Maintenance Job](#) on page **336**

Search

The topics in the following section cover searching for objects, downloading search results, and browsing the Inventory.

Search Overview

Use IBM Spectrum Copy Data Management to explore objects on cataloged providers. With the Search feature, you can easily search for and rapidly find all objects that match certain criteria.

The DellEMC Unity Inventory may include CIFS shares, file systems, hosts, host containers, LUNs, NAS servers, NFS shares, pools, snapshots, storage resources, and storage resource replications.

The IBM Inventory may include FlashCopies, Hosts, IOGroups, MDisks, mirrors, Node Canisters, PortIPs, and volumes.

The NetApp ONTAP Inventory may include aggregates, CIFS shares, files, LUNs, networks, NFS exports, nodes, policies, protocols, qtrees, quotas, SnapMirrors, Snapshots, SnapVaults, SVMs, vFilers, and volumes.

The Recovery Inventory may contain datacenters, data stores, ESX hosts, LUNs, folders, recovery points, vApps, vDisks, vSnapshots, and vSpheres.

The VMware Inventory may include data stores, ESX hosts, LUNs, virtual disks, virtual machines, VMware hosts, and virtual snapshots.

The Pure Storage FlashArray Inventory may include volumes, snapshots, LUNS, hosts, and host groups.

You can match a character pattern and apply other filters such as category, object type, and location through the Search feature. The results are presented in the user interface and are also exportable. Furthermore, click an object that appears in the Search results to open a tab with additional details about that object. You can also review previous versions of your files, along with their Snapshot, SnapVault and SnapMirror replication status.

Alternatively, use the Inventory Browser to browse through the list of providers. Drill into the Inventory Browser to logically view the details of the objects underlying a storage system, virtual host, or application.

Use the Time Machine control to view the Inventory as it appeared on a past date.

WHY IT MATTERS:

IBM Spectrum Copy Data Management lets you quickly and easily locate every version of a file across your entire Enterprise. IBM Spectrum Copy Data Management searches its databases and returns results in moments; every instance and version of a file displays across all devices and snapshots.

RELATED TOPICS:

- [Search for Objects](#) on page **346**
- [View Object Details](#) on page **351**
- [View NetApp ONTAP File Details](#) on page **352**
- [Find and Restore a File](#) on page **354**
- [Download Search Results](#) on page **356**
- [Browse Inventory](#) on page **357**

Search for Objects

Use the search feature to find objects on providers that are cataloged in IBM Spectrum Copy Data Management. Examples of objects are volumes, files, snapshots, qtrees, and virtual machines. You can tailor your search by applying filters.

There are two types of search, basic and advanced.

Basic search searches all text fields. Enter a character pattern including wildcards and inline search strings for more advanced searches. IBM Spectrum Copy Data Management searches the entire Inventory and returns all objects with a name that matches or contains the search entry.

Advanced search is similar to basic search, with an additional function. You can search and filter by object name, category, object type, and location. When searching for NetApp ONTAP files, you can also filter by last modified time, creation time, last accessed time, and file size.

Once you initiate the search and it completes, IBM Spectrum Copy Data Management returns all search results defined by your criteria. Click an object that appears in the Search results to open a tab with additional details about that object.

BEFORE YOU BEGIN:

- You can only search for objects on providers that are registered and cataloged. See [Register a Provider](#) on page 75 and Plan Overview.

To search for objects:

1. Click the **Search**  tab.
2. Open a new Search pane. If this is the first search in your IBM Spectrum Copy Data Management session, click the **Search**  tab. If you have already done a Search, go into an existing Search pane and click New Search. From a new Search pane, you can perform a basic search or an advanced search.

To perform a basic search:

1. In the **Enter search term** field, enter the character pattern to search on. Following are guidelines for entering Search terms:
 - Enter a character string to find objects with a name that matches or contains the character string. You can also enter partial character strings. Character strings are case insensitive.
 - Enter ***** to return all available objects.
 - Apply wildcards as needed. Wildcard considerations are described later in this topic.
2. Click **Search Now**. The list of objects that meet all the criteria displays.
3. Click an object name. The properties of the object display in a new tab. The specific properties vary by type of object.

To perform a basic search using inline search parameters:

Using the following inline search strings, you can perform complex searches based on a file's location, size, and access, creation, or modified time from the basic search field.

Search by object location:

Limit your search to a specific cataloged location using the following examples:

`type:file location:<HOSTNAME>*` searches for all objects on the storage system associated with the entered host

`type:file location:<HOSTNAME>* name:*.txt` searches for .txt files on the storage system associated with the entered host

Search by object size:

Search for cataloged objects with a specific file size or file size range using the following examples:

`size:100KB` searches for all objects that are 100 KB in size

`size:50KB-100MB` searches for all objects between 50 KB and 100 MB in size

`size:*-100MB` searches for all objects that are less than 100 MB in size

`size:100MB-*` searches for all objects that are larger than 100 MB in size

The following size unit strings are supported:

k, K, KB, Kb, kB, kb, KiB, kib, kilobyte, and kilobytes

m, M, MB, Mb, mB, mb, MiB, mib, megabyte, and megabytes

g, G, GB, Gb, gB, gb, GiB, gib, gigabyte, and gigabytes

t, T, TB, Tb, tB, tb, TiB, tib, terabyte, and terabytes

p, P, PB, Pb, pB, pb, PiB, pib, petabyte, and petabytes

Search by object access, creation, and modified time:

Search for cataloged objects that were last accessed, modified, or created at a specific time or time range using the following examples:

`atime:2yearsago` searches for all objects with an access time of two years ago from the time of the search. `ctime` searches against the object's creation time, and `mtime` searches against the object's modification time.

`atime:2yearsago-lastyear` searches for all objects with an access time between last year and two years ago. `ctime` searches against the object's creation time, and `mtime` searches against the object's modification time.

`atime:past2weeks` searches for all objects with an access time from the past two weeks. `ctime` searches against the object's creation time, and `mtime` searches against the object's modification time.

The following time strings are supported:

years, yearsago, year, yearago
 months, monthsago, month, monthago
 weeks, weeksago, week, weekago
 days, daysago, day, dayago
 hours, hoursago, hour, hourago
 minutes, minutesago, minute, minuteago

Combining search strings:

By combining the above search strings in the basic search field, you can limit your search to specific objects, locations, and size ranges.

```
*.vmdk type:file location:<HOSTNAME>/vmtemplates/* catalog:netapp
size:2MB-5MB
```

In this example, search results include all resources that include ".vmdk," residing on a resource named <HOSTNAME>/vmtemplates and its subfolders within a NetApp ONTAP catalog, with a size greater than 2 MB but less than 5 MB.

To perform an advanced search:

1. Click **Advanced Search**.
2. On the **Advanced Search** dialog, enter filters:

Search For

The resource category includes Applications, DellEMC Unity, IBM, NetApp ONTAP, Recovery, or VMware and their associated object types.

Note: Results returned from a low-level NetApp ONTAP file search differ from other object search results. On the searched file's properties pane, you can review previous versions of your files, along with their Snapshot, SnapVault and SnapMirror replication status on the file's properties pane.

For the VMware catalog, object types include Datacenter, Datastore, ESX Host, ESX LUN, Folder, Recovery Point, vApp, vDisk, VM, vSnapshot, and vSphere.

Select **All** for all categories and object types.

Name

Object name or character pattern.

Enter a character string to find objects with a name that matches or contains the character string. You can also enter partial character strings. Character strings are case insensitive.

Enter ***** to return all available objects.

Apply wildcards as needed. Wildcard considerations are described later in this topic.

Location

The place where the object resides. This is usually the host name or the host/volume. Wildcards can be used.

Hide Duplicates

Toggles the behavior of duplicate search results. The default option, **No**, displays duplicate search results in the search results pane. Select **Yes** to hide duplicate search results. View the object's properties to view duplicate versions of an object.

In some cases, the name of a returned object on the search results pane may be the same as another object, however the resources where the objects reside is different. Review the file properties of the objects by selecting their names on the search results pane to view the differences between the returned entries.

The following filters apply to advanced NetApp ONTAP File searches only. Select **NetApp ONTAP > File** in the **Search For** dialog to view the following filters:

Last Modified Time, Creation Time, Last Accessed Time

Filter a search by modification, creation, and accessed dates with the calendar tool. Select **On or after** and **On or before** to set a date range.

File Size

Filter a search by a file size range. Enter a file size and select bytes, kilobytes, megabytes, or gigabytes.

3. Click **Search** . The list of objects that meet all the criteria displays.
4. Click an object name. The properties of the object display in a new tab. The specific properties vary by type of object.

Tip: Periodically closing tabs helps simplify navigation and browsing. To close multiple tabs, right-click a tab then select **Close Tab**, **Close Other Tabs**, or **Close All Tabs**.

Wildcard considerations:

A wildcard is a character that you can substitute for zero or more unspecified characters when searching text. Position wildcards at the beginning, middle, or end of a string, and combine them within a string.

- Match a character string with an asterisk, which represents a variable string of zero or more characters:
 - string*** searches for terms like string, strings, or stringency
 - str*ing** searches for terms like string, straying, or straightening
 - *string** searches for terms like string or shoestring

- Match a single character with a question mark:
 - string?** searches for terms like strings, stringy, or string1
 - st??ring** searches for terms like starring or steering
 - ???string** searches for terms like hamstring or bowstring

You can use multiple asterisk wildcards in a single text string, though this might considerably slow down a large search.

NEXT STEPS:

- You can download the search results as a CSV file format. See [Download Search Results](#) on page **356**.
- You can reorder and resize columns in the search results table.
- You can learn more about an object or its attributes in the search results. See [View Object Details](#) on page **351**.

RELATED TOPICS:

- [Download Search Results](#) on page **356**
- [Browse Inventory](#) on page **357**
- [View Object Details](#) on page **351**
- [View NetApp ONTAP File Details](#) on page **352**
- [Search and Filter Guidelines](#) on page **548**
- [Select, Sort, and Reorder Columns](#) on page **551**

View Object Details

From the search results, you can view the attributes of a searched object, including its location, type, and the dates associated with its creation and modification. If more than one version of an object exists in the Inventory, the search results displays the latest version and attributes of the object.

Note that this topic does not apply to searches for low-level NetApp ONTAP files cataloged with a NetApp ONTAP File Inventory job. Results returned from a NetApp ONTAP file search differ from other object search results.

BEFORE YOU BEGIN:

- You can only browse objects on providers that are registered and cataloged.
- Search for an object. Note that you can filter search results to only return specific objects by using an advanced search.

To view object details:

1. Click an object name in the search results pane to view more information about an object. The properties of the object display in a new tab.
2. If more than one version of the object exists in the Inventory, select a version of the object to view from the Versions tab.

Note: The object that displays in the search results pane is the latest version. In some cases, the name of a returned object on the search results pane may be the same as another object, however the resources where the objects reside is different. Review the file properties of the objects by selecting their names on the search results pane to view the differences between the returned entries.

3. Review the properties of the object. The specific properties vary by type of object.

RELATED TOPICS:

- [View NetApp ONTAP File Details](#) on page **352**
- [Search Overview](#) on page **345**
- [Search for Objects](#) on page **346**

View NetApp ONTAP File Details

From the search results, you can view the attributes of a searched file, including its location, size, and the date the file was added to the Inventory. You can also quickly view a file's creation time, last modified time, and the last accessed time from the search results pane.

Note that this topic only applies to searches for low-level NetApp ONTAP files cataloged with a NetApp ONTAP File Inventory job.

On the searched file's properties pane, you can review previous versions of your files, along with their Snapshot, SnapVault, and SnapMirror replication status on the file's properties pane.

If more than one version of a file exists in the Inventory, the search results pane displays the latest version and attributes of the file. View the file's properties to view previous versions and associated attributes of the file.

Note that you can filter search results to only return files through an advanced search.

WHY IT MATTERS:

IBM Spectrum Copy Data Management provides every location of a file so you can ensure the file is fully protected. You can also quickly identify data leakage or determine if you are wasting valuable resources by overprotecting a file.

BEFORE YOU BEGIN:

- You can only view all available file details once your NetApp ONTAP providers are cataloged through both a NetApp ONTAP Storage Inventory job and a NetApp ONTAP File Inventory job. By enabling the **Catalog all available snapshots** option in the NetApp ONTAP File Inventory job, you can view multiple versions of a file. See [Create an Inventory Job Definition - NetApp ONTAP Storage](#) on page 183 and [Create an Inventory Job Definition - NetApp ONTAP File](#) on page 186.
- Search for a file by selecting **NetApp > File** from the Search For filter on the Advanced Search dialog. See [Search for Objects](#) on page 346.

To view file details:

1. Click a file name in the search results pane to view more information about a file. The properties of the file display in a new tab.
2. If more than one version of the file exists in the Inventory, select a version of the file to view from the Versions tab.

Note: The file displayed in the search results pane is the latest version. In some cases, the name of a returned object on the search results pane may be the same as another object, however the resources where the objects reside is different. Review the file properties of the objects by selecting their names on the search results pane to view the differences between the returned entries.

3. Click an instance in the Snapshots, SnapMirrors, or SnapVaults tab to view details about the associated replication instance, such as the creation date, mirror or vault location, and the source.

NEXT STEPS:

- Restore a file to a previous version. See [Find and Restore a File](#) on page **354**.

RELATED TOPICS:

- [View Object Details](#) on page **351**
- [Search Overview](#) on page **345**
- [Search for Objects](#) on page **346**
- [Find and Restore a File](#) on page **354**

Find and Restore a File

Access a file on a network-connected Windows machine directly from the search results pane, then restore the file to a previous version using snapshot technology. Search for a file and IBM Spectrum Copy Data Management returns the location of every instance and version of that file across your NetApp ONTAP storage infrastructure.

BEFORE YOU BEGIN:

- You can only browse files on providers that are registered and cataloged. Note that increasing the number of catalog instances to keep increases the number of versions that display on a file's properties pane. See [Create an Inventory Job Definition - NetApp ONTAP File](#) on page 186.
- To enable NFS links from the **File Links** tab on an object's property pane, access IBM Spectrum Copy Data Management through a browser in a Linux environment. Additionally, an automounter must be running using the default local path of /net.
- By default, CIFS users cannot see the .snapshot directory. Set the **cifs.show_snapshot** option to **On** to view the .snapshot directory.
- Activate snapshot functionality on the host and enable the necessary access privileges for the user accessing the host. See Use a Script Job to Run Snap Creator.

To find and restore a file:

To restore a file on a Windows machine, search for the file in IBM Spectrum Copy Data Management and discover its location through the file's properties pane. Note that the Windows machine must have network access with snapshot functionality enabled and the volume where the file resides must be shared through CIFS.

1. Search for a file to restore, then click the file name in the search results pane to access the file's properties pane.
2. Select a version of the file to restore from the Versions tab.
3. Click the **File Links** tab.
4. Copy the **CIFS** file path.
5. Paste the path in Windows Explorer. Remove the file name from the path to view the file's containing folder. Press **Enter**.
6. Right-click the file to restore in Windows Explorer and select **Restore previous versions**.
7. Click the **Previous Versions** tab to view previous versions of the file.
8. On the file's properties pane, perform one of the following actions:
 - Select a previous version of the file and click **Open** to open the previous version without altering the original version.

Select a previous version of the file and click **Copy** to copy the previous version to an alternate location.

Select a previous version of the file and click **Restore** to replace or roll back the file to the selected previous version.

RELATED TOPICS:

- [View NetApp ONTAP File Details](#) on page **352**
- [Search Overview](#) on page **345**
- [Search for Objects](#) on page **346**

Download Search Results

After performing a search, use the download feature to save search results as a CSV file. You can view the file or save it for offline viewing.

Note: Hyperlinks might not be active in the CSV file. Only search results that are visible on the search results pane are exported.

BEFORE YOU BEGIN:

- Use the search feature to find objects on providers that are cataloged in IBM Spectrum Copy Data Management. See [Search for Objects](#) on page **346**.

To save search results as a CSV file:

1. Perform a basic or advanced search.
2. In the search results pane, click **Download**.
3. Select Open to view the file now or Save to save the file to your local disk.
4. Click **OK**.

Tip: A CSV file can be opened with Microsoft Excel.

Tip: Times appear as Epoch timestamps, which can be converted using any third party Epoch timestamp converter.

RELATED TOPICS:

- [Search for Objects](#) on page **346**
- [View Object Details](#) on page **351**
- [View NetApp ONTAP File Details](#) on page **352**

Browse Inventory

Use IBM Spectrum Copy Data Management to explore resources that are cataloged and find the properties of underlying objects. View the details for a storage system, virtual host, or application server.

The DellEMC Unity Inventory may include CIFS shares, file systems, hosts, host containers, LUNs, NAS servers, NFS shares, pools, snapshots, storage resources, and storage resource replications.

The IBM Inventory may include FlashCopies, Hosts, IOGroups, MDisks, mirrors, Node Canisters, PortIPs, and volumes.

The NetApp ONTAPInventory may include aggregates, CIFS shares, files, LUNs, networks, NFS exports, nodes, policies, protocols, qtrees, quotas, SnapMirrors, Snapshots, SnapVaults, SVMs, vFilers, and volumes.

The Recovery Inventory may contain datacenters, data stores, ESX hosts, LUNs, folders, recovery points, vApps, vDisks, vSnapshots, and vSpheres.

The VMware Inventory may include data stores, ESX hosts, LUNs, virtual disks, virtual machines, VMware hosts, and virtual snapshots.

The Pure Storage FlashArray Inventory may include volumes, snapshots, LUNS, hosts, and host groups.

BEFORE YOU BEGIN:

- You can only browse objects on providers that are registered and cataloged.

To browse providers that are cataloged:

1. Click the **Search**  tab. On the Views pane, select **Catalog** . The Inventory Browser opens.
2. Drill down through providers in the Inventory Browser. A tab that displays the names of the underlying objects opens.
3. In the tab, click an object name. The properties of the object display in a new tab. The specific properties vary by type of object.

To browse through the Inventory on a previous date:

1. Click the **Search**  tab. On the Views pane, select **Catalog** . The Inventory Browser opens.
2. Click the **All Resources** tab. The Time Machine control appears.
3. Use the calendar tool to set the Time Machine to a previous date. The Inventory Browser reflects the state of the Inventory on that date.
4. Drill down through providers in the Inventory Browser. A tab that displays the names of the underlying objects opens.

Tip: Periodically closing tabs helps simplify navigation and browsing. To close multiple tabs, right-click a tab then select **Close Tab**, **Close Other Tabs**, or **Close All Tabs**.

RELATED TOPICS:

- [Search Overview](#) on page **345**
- [Search for Objects](#) on page **346**
- [View Object Details](#) on page **351**
- [View NetApp ONTAP File Details](#) on page **352**

Report

The topics in the following section cover running and customizing reports, as well as individual report details.

Report Overview

IBM Spectrum Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

The information in these reports are presented in a chart-based **Quick View** section, or tabular **Summary View** and **Detail View** sections.

Reports include interactive elements, such as searching for individual values within a report, vertical scrolling, and column sorting. Information groups, such as the Primary Source Volume groups in the NetApp ONTAP Protection Usage report, can also be sorted by clicking the group name.

You can add a Report job to summarize information about cataloged providers and the data and other resources that reside on them, then schedule the Report job to run as defined by the parameters of the schedule.

To further analyze the data or print a hard copy, use the export functionality to save the data from the generated report to an Adobe PDF, Microsoft Word file, Microsoft Excel file, or HTML file.

RELATED TOPICS:

- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**
- [Edit a Customized Report](#) on page **364**
- [Download a Report](#) on page **365**
- [Create a Report Job Definition](#) on page **337**
- [Application Reports](#) on page **368**
- [System Management Reports](#) on page **376**
- [File Analytics Reports](#) on page **388**
- [Protection Compliance Reports](#) on page **398**
- [Storage Protection Reports](#) on page **435**
- [Storage Utilization Reports](#) on page **449**

Run a Report

Perform the following steps to run any report from the Report tab. You can run reports with predefined default parameters or run customized reports driven by custom parameters.

BEFORE YOU BEGIN:

- At least one provider must be associated with a job definition. See [Register a Provider](#) on page 75.
- Run a job immediately; sometimes referred to as on demand. See [Start, Pause, and Hold a Job Session](#) on page 162.
- You can also create a schedule that can be applied to a job definition so that the job is run as defined by the parameters of the schedule. See [Create a Schedule](#) on page 155.

To run a report using default parameters:

1. Click the **Report**  tab. On the Views pane, select **Reports** . The Report Browser opens.
2. Select a predefined report from the Report Browser pane.
3. Click **Run** . The default report data is returned in the Report pane.

To run a report using custom parameters:

1. Click the **Report**  tab. On the Views pane, select **Reports** . The Report Browser opens.
2. Select a predefined report from the Report Browser pane.
3. Select report parameter values in the Parameters pane. Parameters are unique to each report. The parameters that you select drive the report output.
4. Click **Run** . The customized report data is returned in the Report pane.

NEXT STEPS:

- Create a Report job definition to schedule the report to run automatically. See [Create a Report Job Definition](#) on page 337.
- Save a report with customized parameters. See [Create a Customized Report](#) on page 363.

RELATED TOPICS:

- [Create a Customized Report](#) on page 363
- [Create a Report Job Definition](#) on page 337

- [Download a Report](#) on page **365**
- [Report Overview](#) on page **360**

Create a Customized Report

Perform the following steps to create a report with customized parameters. Select a predefined report, set custom parameters, and save the report with a customized name to run on demand or create a schedule to run the report as defined by the parameters of the schedule.

Customized reports display nested under their predefined source report on the Report Browser pane.

BEFORE YOU BEGIN:

- At least one provider must be associated with a job definition. See [Register a Provider](#) on page 75.

To create a customized report:

1. Click the **Report**  tab. On the Views pane, select **Reports** . The Report Browser opens.
2. Select a predefined report to save as a customized report from the Report Browser pane.
3. Select report parameter values in the Parameters pane. Parameters are unique to each report. The parameters that you select drive the report output.
4. Click **Save As** . The Save As window opens.
5. Enter a **Title** and a **Description** for the customized report. Report names can include alphanumeric characters and the following symbols: \$ - _ . + ! * ' ().
6. Click **Submit**. The customized report is saved.
7. Return to the Report Browser. Expand the original predefined report to view associated customized reports.

NEXT STEPS:

- Run the customized report from the **Report**  tab. See [Run a Report](#) on page 361.

RELATED TOPICS:

- [Edit a Customized Report](#) on page 364
- [Run a Report](#) on page 361
- [Download a Report](#) on page 365
- [Report Overview](#) on page 360

Edit a Customized Report

Perform the following steps to edit a report with customized parameters. Edit the custom parameters, and save the report with a customized name to run on demand or create a schedule to run the report as defined by the parameters of the schedule.

Customized reports display nested under their predefined source report on the Report Browser pane.

BEFORE YOU BEGIN:

- At least one provider must be associated with a job definition. See [Register a Provider](#) on page 75.
- Run a job immediately; sometimes referred to as on demand. See [Start, Pause, and Hold a Job Session](#) on page 162.
- Create a customized report. See [Create a Customized Report](#) on page 363.

To edit a customized report:

1. Click the **Report**  tab. On the Views pane, select **Reports** . The Report Browser opens.
2. Select a customized report from the Report Browser pane.
3. Edit report parameter values in the Parameters pane.
4. Click **Save** . The customized report is saved.
5. Return to the Report Browser. Expand the original predefined report to view associated customized reports.

NEXT STEPS:

- Run the customized report from the **Report**  tab. See [Run a Report](#) on page 361.

RELATED TOPICS:

- [Create a Customized Report](#) on page 363
- [Run a Report](#) on page 361
- [Download a Report](#) on page 365
- [Report Overview](#) on page 360

Download a Report

Download reports from the report output or from the Jobs History pane. Reports can be downloaded as HTML files, Adobe PDFs, Microsoft Excel spreadsheets, and Microsoft Word files.

Best Practice: Generated reports are automatically removed from the **Jobs History**  pane seven days after their initial run. To save a report indefinitely, download it as an HTML file, Adobe PDF, Microsoft Excel spreadsheet, or Microsoft Word file. Click **Download**  while viewing an open report or download a previously generated report from the **Jobs History**  pane.

Note: Some report images appear truncated when viewed in Microsoft Word. For best results, view downloaded reports through the Web Layout.

BEFORE YOU BEGIN:

- Run a report from the Report Browser pane. See [Run a Report](#) on page 361.

To download a report:

1. Run a report from the Report Browser pane.
2. Click **Download** .
3. Select **HTML** , **PDF** , **Excel** , or **Word** . You can view the report now or save it to a file.

To download a generated report:

1. Click the **Report**  tab. On the Views pane, select **Jobs History** .
2. On the All Reports pane, select a generated report to download.
3. Click the **HTML** , **PDF** , **Excel** , or **Word**  icon associated with the report. You can view the report now or save it to a file.

Reports are saved on the IBM Spectrum Copy Data Management appliance in the following directories:

- Report Download location: `/data/reports/output`
- Report Name location: `/data/reports/repository/reports/output`
- Report Temporary location: `/data/reports/repository/reports/doc`

NEXT STEPS:

- Review your downloaded report for further analysis of your data.

RELATED TOPICS:

- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Delete a Generated Report

Delete generated reports from the list of generated reports on the Jobs History pane.

Best Practice: Generated reports are automatically removed from the **Jobs History**  pane seven days after their initial run. To save a report indefinitely, download it as an HTML file, Adobe PDF, Microsoft Excel spreadsheet, or Microsoft Word file. Click **Download**  while viewing an open report or download a previously generated report from the **Jobs History**  pane.

To delete a generated report:

1. Click the **Report**  tab. On the Views pane, select **Jobs History** .
2. On the All Reports pane, select a generated report to delete.
3. Click **Delete** . A confirmation dialog box opens.
4. Click **Yes** to delete the generated report.

RELATED TOPICS:

- [Download a Report](#) on page **365**
- [Run a Report](#) on page **361**

Application Reports

The Application Reports help you review your application server configuration. Use the Application Reports to view your application server's database, log disks, and eligibility for protection. Reports are based on the data collected by the most recently run job.

Choose the Application reports that fit your needs:

- **Application Configuration Report** - Review the configuration of your application servers including the disks you the database and logs reside.
- **Application RPO Compliance Report** - The Application RPO Compliance report displays your application servers in relation to your recovery point objective parameters.

RELATED TOPICS:

- [Application Configuration Report](#) on page **369**
- [Application RPO Compliance Report](#) on page **372**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**
- [Download a Report](#) on page **365**

Application Configuration Report

Review the configuration of your application servers including the disks where the database and logs reside.

BEFORE YOU BEGIN:

- Create and run a Database Inventory job. You can select one or more application provider in a single job definition for cataloging. See [Create an Inventory Job Definition - Database](#) on page 170.
- For Microsoft-based application servers, ensure the vCenter containing the virtual machines are registered in IBM Spectrum Copy Data Management. Ensure that the virtual machine credentials for the virtual machines hosting the applications are configured. See [Add a Credential](#) on page 132.

Parameters

Use the following parameters to customize your report:

- Application Type: File System, InterSystems Caché , Microsoft SQL, Oracle, SAP HANA. Multiple selections are supported.

The default report parameter, **All**, reports on all available application types in your configuration.

Detail View - Microsoft SQL

The following fields and corresponding data display in the Microsoft Active Directory, Microsoft Exchange, Microsoft SQL, and Microsoft SharePoint sections of the Application Configuration report:

Instance

The name and location of the SQL instance.

Database

The name of the database associated with the SQL instance and application server.

Data Disk(s)

The name of the associated database disk mount points, along with the associated volume and storage array in parentheses.

Log Disk(s)

The name of the associated log disk mount points, along with the associated volume and storage array in parentheses.

Eligible for Protection

The status of the database's eligibility for protection.

Detail View - Oracle

The following fields and corresponding data display in the Oracle section of the Application Configuration report:

Oracle Home

The name and location of the Oracle home.

Database

The name of the database associated with the Oracle home and application server.

Data Disk(s)

The name of the associated database disk mount points, along with the associated volume and storage array in parentheses.

Log Disk(s)

The name of the associated log disk mount points, along with the associated volume and storage array in parentheses.

Eligible for Protection

The status of the database's eligibility for protection.

Detail View - InterSystems Caché

The following fields and corresponding data display in the InterSystems Caché section of the Application Configuration report:

Instance

The name and location of the InterSystems Caché instance.

Database

The name of the database associated with the InterSystems Caché instance and application server.

Data Disk(s)

The name of the associated database disk mount points, along with the associated volume and storage array in parentheses.

Log Disk(s)

The name of the associated log disk mount points, along with the associated volume and storage array in parentheses.

Eligible for Protection

The status of the database's eligibility for protection.

Detail View - SAP HANA

The following fields and corresponding data display in the SAP HANA section of the Application Configuration report:

Instance

The name and location of the SAP HANA instance.

Database

The name of the database associated with the SAP HANA instance and application server.

Data Disk(s)

The name of the associated database disk mount points, along with the associated volume and storage array in parentheses.

Log Disk(s)

The name of the associated log disk mount points, along with the associated volume and storage array in parentheses.

Eligible for Protection

The status of the database's eligibility for protection.

Detail View - File System

The following fields and corresponding data display in the File System section of the Application Configuration report:

File System

The name of the database associated with the file system.

Mount Points

The name of the associated database disk mount points, along with the associated volume and storage array in parentheses.

File System Type

The type of registered file system.

Eligible for Protection

The status of the database's eligibility for protection.

RELATED TOPICS:

- [Application Reports](#) on page **368**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Application RPO Compliance Report

The Application RPO Compliance report displays your application servers in relation to your recovery point objective parameters. Determine which of your application servers are not in compliance with your RPO parameters, and discover the reasons for their non-compliance.

BEFORE YOU BEGIN:

- Create and run a Database Inventory job. You can select one or more application provider in a single job definition for cataloging. See [Create an Inventory Job Definition - Database](#) on page 170.
- Create and run a Database Backup job.

Parameters

Use the following parameters to customize your report:

- Application Type
InterSystems Caché , Microsoft SQL, Oracle, SAP HANA. Multiple selections are supported.
- Application Server
Multiple selections are supported.
- Protection Type
Set the protection types to return in the report. Values include Primary and Replication. Multiple selects are supported.
- Storage Vendor
Set the storage vendor types to display in the report. Multiple selections are supported.
- Display Databases That Are
Set the compliance status of your application server to return in the report. Values include Compliant and Not Compliant. Multiple selections are supported. By default, this parameter is set to Not Compliant.
- RPO Older Than
Set the age of the recovery point objective in days.

The default report parameters report on all non-compliant application servers based on an RPO older than one day.

Quick View

The Quick View section displays a bar graph of compliant and non-compliant application servers based on your RPO parameters.

Not Compliant for Primary Protection

The following fields and corresponding data display in the Not Compliant for Primary Protection section of the Application RPO Compliance report:

Database

The names and properties of the resource.

Application Server (Type)

The application server's location and type.

Application Instance

The name of the application instance, or Oracle home if reporting on Oracle databases.

Job Name

The Backup job associated with the application server.

Storage Vendor

The storage vendor associated with the application server.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Reason

The reason the application server does not meet your RPO compliance parameters. Examples include no successful runs of a backup job, or backing up to an unsupported disk.

Compliant for Primary Protection

The following fields and corresponding data display in the Compliant for Primary Protection section of the Application RPO Compliance report:

Database

The names and properties of the resource.

Application Server (Type)

The application server's location and type.

Application Instance

The name of the application instance, or Oracle home if reporting on Oracle databases.

Job Name

The Backup job associated with the application server.

Storage Vendor

The storage vendor associated with the application server.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your application server will be non-compliant.

Not Compliant for Replication

The following fields and corresponding data display in the Not Compliant for Replication section of the Application RPO Compliance report:

Database

The names and properties of the resource.

Source

The source of the secondary protection.

Destination

The destination of the secondary protection.

Job Name

The Backup job associated with the application server.

Storage Vendor

The storage vendor associated with the application server.

Protection Time

The most recent instance of a successful run of the Backup job.

Reason

The reason the application server does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Replication

The following fields and corresponding data display in the Compliant for Replication section of the Application RPO Compliance report:

Database

The names and properties of the resource.

Source

The source of the secondary protection.

Destination

The destination of the secondary protection.

Job Name

The Backup job associated with the application server.

Storage Vendor

The storage vendor associated with the application server.

Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your application server will be non-compliant.

RELATED TOPICS:

- [Application Reports](#) on page **368**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

System Management Reports

System Management Reports offer an in-depth view of the status of your IBM Spectrum Copy Data Management configuration, including cataloged storage system information, jobs, and their status.

Use the System Management Reports to answer questions such as:

- What is the operating system, memory, and processor information of my storage system?
- How many jobs are associated with an Inventory?
- What is the average runtime of a job, and when was it last successfully run?

Choose the System Management report that fits your needs:

- **Catalog Summary Report** - Displays a summary of all the high level objects cataloged.
- **Configuration Report** - Review IBM Spectrum Copy Data Management node configuration and associated jobs.
- **Job Report** - Review jobs and their status, including the last time they ran successfully.
- **Job Sessions Report** - Review sessions for jobs and their status.

RELATED TOPICS:

- [Catalog Summary Report](#) on page **377**
- [Configuration Report](#) on page **379**
- [Job Report](#) on page **383**
- [Job Sessions Report](#) on page **386**
- System Sizing Report
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**
- [Download a Report](#) on page **365**

Catalog Summary Report

Review a summary of your cataloged high level objects. Run the Catalog Summary Report to display information about NetApp ONTAP, Application, and VMware objects and the number of associated records cataloged.

Use the Catalog Summary report to answer questions such as:

- What are the total number of files across all cataloged high level objects?
- What is the size of an object's Catalog?
- What are the total number of snapshots on a NetApp ONTAP object?

Parameters

Use the following parameters to customize your report:

- Catalog Type: All, Application, DellEMC Unity, IBM Spectrum Accelerate, IBM Spectrum Virtualize, NetApp ONTAP, NetApp ONTAP Node Summary, Pure Storage FlashArray, and VMware. Multiple selections are supported.

The default report parameter, **All**, reports on all available Catalog information.

NetApp ONTAP Node Summary

The NetApp ONTAP Node Summary section of the report displays an overview of your NetApp ONTAP nodes, including the number of aggregates, volumes, snapshots, and the total number of files on your nodes.

NetApp ONTAP Objects, VMware Objects, Application Objects, IBM Spectrum Virtualize Objects, IBM Spectrum Accelerate Objects, DellEMC Unity Objects, and Pure Storage FlashArray Objects

The following fields and corresponding data display in the NetApp ONTAP Objects, VMware Objects, Application Objects, and IBM Objects, DellEMC Unity Objects, and Pure Storage FlashArray Objects sections of the Catalog Summary report:

Object

The name of the cataloged high level object.

Total Count

The number of records found on the cataloged object.

Catalog Size

The total amount of data on the cataloged object.

Count from Last Update

The latest number of records cataloged during the last update.

Last Updated

The date and time of the most recent Catalog update.

RELATED TOPICS:

- [System Management Reports](#) on page **376**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Configuration Report

Review the IBM Spectrum Copy Data Management providers added to IBM Spectrum Copy Data Management and their associated configurations.

Use the Configuration report to answer questions such as:

- What is the operating system, memory, and processor information of my nodes?

Parameters

Use the following parameters to customize your report:

- Configuration Type: All, AWS Nodes, Application Nodes, DellEMC Unity Nodes, IBM Spectrum Accelerate Nodes, IBM Spectrum Protect Snapshot Nodes, IBM Spectrum Virtualize Nodes, NetApp ONTAP Nodes, Pure Storage FlashArray Nodes, and VMware Nodes. Multiple selections are supported.

The default report parameter, **All**, reports on all available nodes in your configuration.

NetApp ONTAP Nodes

The following fields and corresponding data display in the NetApp ONTAP Nodes section of the Configuration report:

Node (Site)

The name of the NetApp ONTAP node and the associated site.

Model

The model number of the NetApp ONTAP node.

Type

The node type.

Operating System Version

The installed operating system version on the NetApp ONTAP node.

Memory

The amount of memory installed on the NetApp ONTAP node.

Processor(s)

The number of processors on the NetApp ONTAP node.

VMware Nodes

The following fields and corresponding data display in the VMware Nodes section of the Configuration report:

Node (Site)

The name of the vCenter node and the associated site.

Type

The node type.

Product Version

The installed product version on the vCenter node.

OS Type

The installed operating system type on the vCenter node.

IBM Spectrum Virtualize Nodes

The following fields and corresponding data display in the IBM Spectrum Virtualize Nodes section of the Configuration report:

Node (Site)

The name of the IBM Spectrum Virtualize node and the associated site.

Model

The Spectrum Virtualize storage model name.

Operating System Version

The installed operating system version on the IBM Spectrum Virtualize node.

IBM Spectrum Protect Snapshot Nodes

The following fields and corresponding data display in the IBM Spectrum Protect Snapshot Nodes section of the Configuration report:

Node (Site)

The name of the IBM Spectrum Protect Snapshot node and the associated site.

IBM Spectrum Protect Snapshot Version

The IBM Spectrum Protect Snapshot version number.

vCenter Server (User)

The associated vCenter server and user.

IBM Spectrum Accelerate Nodes

The following fields and corresponding data display in the IBM Spectrum Accelerate Nodes section of the Configuration report:

Node (Site)

The name of the IBM Spectrum Accelerate node and the associated site.

IBM Spectrum Accelerate Version

The IBM Spectrum Accelerate version number.

System Name

The system name of the IBM Spectrum Accelerate node.

DellEMC Unity Nodes

The following fields and corresponding data display in the DellEMC Unity Nodes section of the Configuration report:

Node (Site)

The name of the DellEMC Unity node and the associated site.

Version

The DellEMC Unity version number.

Application Nodes

The following fields and corresponding data display in the Application Nodes section of the Configuration report:

Node (Site)

The name of the Oracle node and the associated site.

Application Type

The type of application server, including SQL or Oracle.

OS Type

The installed operating system type on the application node.

Server Type

The application server type, including physical or virtual.

AWS Nodes

The following fields and corresponding data display in the AWS Nodes section of the Configuration report:

Node (Site)

The name of the AWS node and the associated site.

Region

The AWS node's associated region.

Pure Storage FlashArray Nodes

The following fields and corresponding data display in the Pure Storage FlashArray Nodes section of the Configuration report:

Node (Site)

The name of the Pure Storage FlashArray node and the associated site.

Version

The Pure Storage FlashArray version number.

RELATED TOPICS:

- [System Management Reports](#) on page **376**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Job Report

Review the available jobs in your IBM Spectrum Copy Data Management configuration. Run the Job report to view jobs by type, their average runtime, and their successful run percentage.

Use the Job report to answer questions such as:

- What types of jobs are available?
- What is the average runtime of a job, and when was it last successfully run?
- How many times has a specific job run successfully or failed?

Parameters

Use the following parameters to edit your report:

- Job Type
Multiple selections are supported.
- Days Since Successful Run
- Show Job Details
Enable to display the Job ID field in the Detail View section.

The default report parameters report on all available jobs.

Quick View

The Quick View section displays a pie chart of the number of times a job type successfully completed, failed, or was marked with any of the following statuses: unknown, waiting, running, stopped, partial, skipped, aborted, or stopping. Use the Job Type parameter to display File, NetApp ONTAP, Report, Script, VMware, or all job types.

Note: The Quick View section is only modified through the Job Type parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the Jobs report:

Job Type

The type of job. For example, a catalog or report job.

Jobs

The number of jobs associated with the job type.

Runs

The number of times a job of this type ran.

Completed

The number of times a job of this type successfully completed.

Failed

The number of times a job of this type failed.

Other

The number of times a job of this type was marked with any of the following statuses: unknown, waiting, running, stopped, partial, skipped, aborted, or stopping.

Detail View

The following fields and corresponding data display in the Detail View section of the Jobs report:

Job

The job name.

Job ID

The job ID of the displayed job. This field is controlled by the Show Job Details parameter.

Type

The type of job. For example, a catalog or report job.

Runs

The number of times the job ran.

Completed

The number of times the job successfully completed.

Failed

The number of times the job failed.

Other

The number of times the job was marked with any of the following statuses: unknown, waiting, running, stopped, partial, skipped, aborted, or stopping.

Last Successful Run

The date and time the job last ran successfully.

Average Runtime (Days hh:mm:ss)

The average time it takes to run the job, listed in days, hours, minutes, and seconds.

Success %

The percentage of times the job successfully completed.

RELATED TOPICS:

- [System Management Reports](#) on page **376**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Job Sessions Report

Review sessions for jobs and their status in your IBM Spectrum Copy Data Management environment. Run the Job Sessions report to view job sessions by type, their status, start time, finish time, and duration.

Use the Job Sessions report to answer questions such as:

- What job sessions have run in the environment?
- What are the types of job sessions that have run in the environment?
- What is the status of those job sessions?
- What was the start and end time for each job session?
- What was the duration from start to finish for each job session?

Parameters

Use the following parameters to edit your report:

- Job Type
Multiple selections are supported.
- Job
Multiple selections are supported.
- Status
Multiple selections are supported.
- Protection Type
Multiple selections are supported.
- Job Sessions for Past Number of Days
Enter a number for the number of days.

The default report parameters report on all available jobs and 1 for Job Sessions for Past Number of Days.

Detail View

The following fields and corresponding data display in the Detail View section of the Job Sessions report:

Job

The job name.

Type

The type of job. For example, a catalog or report job.

Status

The status of the job session.

Start Time

The time that the job session started in UTC time.

Finish Time

The time that the job session finished in UTC time.

Duration (Days hh:mm:ss)

The total time that it takes for the job session to run, listed in days, hours, minutes, and seconds.

RELATED TOPICS:

- [System Management Reports](#) on page **376**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

File Analytics Reports

The File Analytics Reports help you review your storage needs and examine your storage capacity. There are a number of variables that you can use to view the files on your system. Using the File Analytics Reports you can drill down to information about the size of files, the age of files, and type of files on your storage systems. Reports are based on the data collected by the most recently run job.

Use the File Analytics Reports to answer questions such as:

- What is the average size of files stored?
- What files have not been accessed within a given time period?
- What is the amount of space usage by file type?

Choose the File Analytics report that fits your needs:

- **File Usage by Owner Report** - Identify the largest space consumers on your storage systems by owner to help manage storage utilization. Run the File Usage by Owner report to view the owners consuming the largest amount of space on your storage systems.
- **Files By Age Report** - Review the age of files on your NetApp storage systems based on the creation date, the last time accessed, and the last time modified.
- **Files By Category Report** - Identify the application types that monopolize your storage. View storage consumption by extension and the number of files associated with the extension.
- **Files By Size Report** - Discover the largest space consumers on your NetApp storage systems. View the largest files, path and volume location, and last time accessed.

Quick View

This area of the report is a graphical illustration of the report using pie charts. For example, the quick view of the Files by Age report shows the age of all files on a volume.

Detail View

This area of the report is a table where each row details a node, its corresponding volume, and details returned by the report. For example, the Files By Size report shows the largest files on your node, their size, and the last time they were accessed.

WHY IT MATTERS:

IBM Spectrum Copy Data Management uses powerful file reporting and analytics to quickly pinpoint inefficient use of storage. You can identify unwanted objects such as old files that have not been accessed for a set amount of time, extremely large files, and file types that may violate storage policies, like video and music files.

RELATED TOPICS:

- [File Usage by Owner Report](#) on page 390
- [Files By Age Report](#) on page 392

- [Files By Category Report](#) on page **394**
- [Files By Size Report](#) on page **396**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**
- [Download a Report](#) on page **365**

File Usage by Owner Report

Identify the largest space consumers on your storage systems by owner to help manage storage utilization. Run the File Usage by Owner report to view the owners consuming the largest amount of space on your storage systems as well as the number of files associated with an owner.

Use the File Usage by Owner report to answer questions such as:

- Which owners are consuming the largest amount of space on a selected storage system?
- How many files are associated with a specific owner on a selected storage system?

BEFORE YOU BEGIN:

- Create and run a NetApp ONTAP File Inventory job with the Traversal Mode set to Filewalk in conjunction with the IBM Spectrum Copy Data Management Filewalker tool. See [Create an Inventory Job Definition - NetApp ONTAP File](#) on page 186.

Parameters

Use the following parameters to customize your report:

- NetApp ONTAP Storage
- Volume
- Limit No. of Owners to View
Select Yes to limit the number of owners to view through the Enter No. of Owners to View parameter.
Select No to display all owners in the report.
- Number of Owners to View
If the Limit No. of Owners to View parameter is set to Yes, set the number of owners to display in the Details section of the report.
- Export Date Format
Set the date format to use when exporting data.

The default report parameters report on owners consuming the most amount of space on all storage systems and volumes.

Quick View

The Quick View section displays a graph of the top ten owners consuming the most amount of space as well as the top ten owners with the largest number of files on your storage systems.

Note: The Quick View section is modified through the NetApp ONTAP Storage, Volume, and Number of Owners to View parameters.

Detail View

The following fields and corresponding data display in the Detail View section of the File Usage by Owner report:

Owner

The assigned volume owner that is consuming the largest amount of space on your selected storage systems.

Total Used

The amount of space on a storage system used by the owner.

File Count

The number of files on the storage system per owner.

RELATED TOPICS:

- [File Analytics Reports](#) on page **388**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Files By Age Report

Review the age of files on your node based on the creation date and the last time modified. Run the Files By Age report to view the age of files on your node, from older than ten year to less than one year.

Age categories include:

- Older than ten years
- Five to ten years
- Four to five years
- Three to four years
- Two to three years
- One to two years
- 180 days to one year
- Less than 180 days

Use the Files by Age report to answer questions such as:

- How much storage space and the percentage of the total space is used by files created more than ten years ago?
- How much storage space is used within a given time period?

BEFORE YOU BEGIN:

- To view accessed time information, create and run a NetApp ONTAP File Inventory job with the Traversal Mode set to Filewalk in conjunction with the IBM Spectrum Copy Data Management Filewalker tool. See [Create an Inventory Job Definition - NetApp ONTAP File](#) on page 186.

Parameters

Use the following parameters to customize your report:

- NetApp ONTAP Storage
- Volume
 - Multiple selections are supported.
- Date accessed, created, or modified.
- File Size Equal To or Greater Than (GB)
 - Set the minimum file size to display. By default, only files greater than 1 GB display.
- Limit to File Extensions

Set the file types to include in the report by their extension. Separate multiple extensions by commas.

- **Export Data**

Set the age of data to export, from greater than 180 days to greater than 10 years. A list of files based on age is exported. Each row in the file contains directory path, file name, type, and other metadata related to the file. By default, this parameter is set to No. Set the location of the export file on the Run Report dialog, which displays before the report is run. By default, the file is exported to the local /data/reports directory.

- **Export Date Format**

Set the date format to use when exporting data.

The default report parameters report on the age of files greater than 1 GB on all storage systems and volumes.

Quick View

The Quick View section displays a pie chart of the storage used by files based on creation date and the last time modified. Use the NetApp ONTAP Storage parameter to display volumes on all storage systems or a specific storage system.

Note: The Quick View section is modified through the Date, Include Deleted Files, NetApp ONTAP Storage, and Volume parameters.

Detail View

The following fields and corresponding data display in the Detail View section of the Files by Age report:

File Age (days)

Age groups including older than ten years, five to ten years, four to five years, three to four years, two to three years, one to two years, one year to 180 days, and less than 180 days.

Total Data

The amount of space used by files in the associated age group.

File Count

The number of files in the associated age group.

RELATED TOPICS:

- [File Analytics Reports](#) on page **388**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Files By Category Report

Identify application types that monopolize your storage. Run the Files By Category report to view files on selected nodes sorted by extension. View the amount of space a file type is using on a node, as well as the number of files associated with the extension.

Use the Files By Category report to answer questions such as:

- What is the amount of space usage by file type?
- What is the count per extension?

Parameters

Use the following parameters to customize your report:

- NetApp ONTAP Storage
- Volume
- Number of Categories to View
- Exclude File Extensions

You can exclude file extensions from the report using the **Exclude File Extensions** parameter. Enter extensions without a leading period, and separate multiple extensions to exclude with a comma. For example, enter **raw, exe, bmp** to exclude files with .RAW, .EXE, and .BMP extensions from the report. Enter a space to exclude files without extensions. Note that filters are not case sensitive.

Note: Files without extensions are grouped in a category labeled {Empty} in the Quick View and Detail View sections.

- Export Date Format

Set the date format to use when exporting data.

The default report parameters report on ten file categories on all storage systems and volumes. Select up to 100 categories.

Quick View

The Quick View section displays graphs of the storage used by a file category and the number of files in a file category. Use the NetApp ONTAP Storage parameter to display volumes on all storage systems or a specific storage system.

Note: The Quick View section is modified through the Number of Categories to View, Exclude File Extensions, NetApp ONTAP Storage, and Volume parameters.

Detail View

The following fields and corresponding data display in the Detail View section of the Files By Category report:

File Extension

The file type included in the report (for example .txt, .exe, or .bin).

Total Used

The amount of space on a storage system used by the file type.

File Count

The number of files associated with the file type.

RELATED TOPICS:

- [File Analytics Reports](#) on page **388**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Files By Size Report

Identify the largest space consumers on your storage systems to help manage storage utilization. Run the Files By Size report to view the largest files, the path and volume where they are located, and their size. View the amount of space the largest file on your storage system utilizes as a percentage of storage system capacity.

Use the Files by Size report to answer questions such as:

- What are the largest files on the system?
- What are the sizes of those files?
- What storage system do they reside on and its associated volume?

BEFORE YOU BEGIN:

- To view file ownership information, create and run a NetApp ONTAP File Inventory job with the Traversal Mode set to Filewalk in conjunction with the IBM Spectrum Copy Data Management Filewalker tool. See [Create an Inventory Job Definition - NetApp ONTAP File](#) on page 186.

Parameters

Use the following parameters to customize your report:

- NetApp ONTAP Storage
- Volume
- Number of Largest Files to View
- Export Date Format

Set the date format to use when exporting data.

The default report parameters report on the ten largest files on all storage systems and volumes. Select up to 100 files.

Quick View

The Quick View section displays a pie chart of the largest files on your storage systems compared to the total storage. Use the NetApp ONTAP Storage parameter to display volumes on all storage systems or a specific storage system.

Note: The Quick View section is modified through the Number of Largest Files to View, NetApp ONTAP Storage, and Volume parameters.

Detail View

The following fields and corresponding data display in the Detail View section of the Files by Size report:

Node

The physical server where your files are stored.

Volume

The name of the volume on the node.

File

The name of the file returned by the report, including the path.

Owner

The assigned volume owner or node where the volume resides. To view file ownership information, create and run a NetApp ONTAP File Inventory job with the Traversal Mode set to Filewalk in conjunction with the IBM Spectrum Copy Data Management Filewalker tool before running this report

No. of Copies

The number of copies of the file that exist in the Inventory.

SnapMirror

The SnapMirror associated with the file.

SnapVault

The SnapVault associated with the file.

Size

The amount of space on the node used by the file.

Last Time Modified

The date and time in which the file was last modified.

RELATED TOPICS:

- [File Analytics Reports](#) on page **388**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Protection Compliance Reports

The Protection Compliance reports help ensure your data is protected through user-defined recovery point objective parameters.

Reports are based on the data collected by the most recently run job. You must catalog all volumes on a storage system using full storage system cataloging to view the correct storage protection status for all volumes and qtrees.

Use the Protection Compliance reports to answer questions such as:

- Which of my NetApp ONTAP or VMware protection jobs have never run successfully?
- What is the remaining compliance time of a specific NetApp ONTAP storage system or VMware object?

Choose the Protection Compliance report that fits your needs.

- **DellEMC Unity RPO Report** - Displays DellEMC Unity storage systems in relation to your recovery point objective parameters.
- **File System RPO Compliance Report** - Displays the RPO compliance of file systems.
- **HPE Nimble Storage RPO Compliance** - Displays HPE Nimble Storage storage systems in relation to your recovery point objective parameters.
- **IBM Spectrum Accelerate RPO Compliance Report** - Displays IBM Spectrum Accelerate storage systems in relation to your recovery point objective parameters.
- **IBM Spectrum Virtualize RPO Compliance Report** - Displays IBM Spectrum Virtualize storage systems in relation to your recovery point objective parameters
- **NetApp ONTAP Protection Usage Report** - Displays the storage usage of IBM Spectrum Copy Data Management protection jobs on the volumes of your NetApp ONTAP storage systems.
- **NetApp ONTAP RPO Compliance Report** - Displays the primary snapshot protection for NetApp ONTAP storage systems.
- **Pure Storage FlashArray RPO Compliance** - Displays Pure Storage FlashArray storage systems in relation to your recovery point objective parameters.
- **Recovery Points Report** - Displays recovery points of various protected resources.
- **Unprotected Virtual Machines Report** - Displays the virtual machines that are not protected as a part of primary snapshot protection.
- **VMware RPO Compliance Report** - Displays the primary snapshot protection for virtual machines and datastores.

RELATED TOPICS:

- [DellEMC Unity RPO Compliance Report](#) on page **400**
- [File System RPO Compliance Report](#) on page **404**
- [HPE Nimble Storage RPO Compliance Report](#) on page **407**
- [IBM Spectrum Accelerate RPO Compliance Report](#) on page **411**
- [IBM Spectrum Virtualize RPO Compliance Report](#) on page **413**
- [NetApp ONTAP Protection Usage Report](#) on page **417**
- [NetApp ONTAP RPO Compliance Report](#) on page **420**
- [Pure Storage FlashArray RPO Compliance Report](#) on page **424**
- [Recovery Points Report](#) on page **427**
- [Unprotected Virtual Machines Report](#) on page **429**
- [VMware RPO Compliance Report](#) on page **431**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**
- [Download a Report](#) on page **365**

DellEMC Unity RPO Compliance Report

The DellEMC Unity RPO Compliance report displays DellEMC Unity storage systems in relation to your recovery point objective parameters. Determine which of your DellEMC Unity storage systems are not in compliance with your RPO parameters, and discover the reasons for their non-compliance.

BEFORE YOU BEGIN:

- Create and run a DellEMC Unity Inventory job. You can select one or more DellEMC Unity provider in a single job for cataloging. See [Create an Inventory Job Definition - DellEMC Unity](#) on page **175**.
- Create and run a DellEMC Unity Backup job. See [Create a Backup Job Definition - DellEMC Unity](#) on page **220**.

Use the DellEMC Unity RPO Compliance report to answer questions such as:

- Which of my DellEMC Unity storage systems are not RPO compliant for replication or snapshot protection?
- Which of my DellEMC Unity Backup jobs have never run successfully?

Parameters

Use the following parameters to customize your report:

- Storage Array
Multiple selections are supported.
- Resource Type
Set the resource type to return in the report. Resource types include file system and LUN.
- Protection Type
Set the RPO compliance protection type. Protection types include replication and snapshot. Multiple selections are supported. By default, this parameter is set to All.
- Display Resources That Are
Set the compliance status of your DellEMC Unity storage systems to return in the report. Values include Compliant and Not Compliant. Multiple selections are supported. By default, this parameter is set to Not Compliant.
- RPO Older Than
Set the age of the recovery point objective in days.

The default report parameters report on all non-compliant DellEMC Unity storage systems based on an RPO older than one day.

Quick View

The Quick View section displays a bar graph of compliant and non-compliant DellEMC Unity storage systems based on your RPO parameters.

Not Compliant for Snapshot Protection

The following fields and corresponding data display in the Not Compliant for Snapshot Protection section of the DellEMC Unity RPO Compliance report:

Resource

The name of the DellEMC Unity storage system.

Resource Type

The type associated with the DellEMC Unity storage system. Resource types include file system and LUN.

Location

The location of the DellEMC Unity storage system.

Job Name

The Backup job associated with the DellEMC Unity storage system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Reason

The reason the DellEMC Unity storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Snapshot Protection

The following fields and corresponding data display in the Compliant for Snapshot Protection section of the DellEMC Unity RPO Compliance report:

Resource

The name of the DellEMC Unity storage system.

Resource Type

The type associated with the DellEMC Unity storage system. Resource types include file system and LUN.

Location

The location of the DellEMC Unity storage system.

Job Name

The Backup job associated with the DellEMC Unity storage system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your DellEMC Unity storage system will be non-compliant.

Not Compliant for Replication

The following fields and corresponding data display in the Not Compliant for Replication section of the DellEMC Unity RPO Compliance report:

Resource

The name of the primary source volume.

Resource Type

The type associated with the DellEMC Unity storage system. Resource types include file system and LUN.

Source

The source of the replication.

Destination

The destination of the replication.

Job Name

The Replication Backup job associated with the DellEMC Unity storage system.

Protection Time

The most recent instance of a successful run of the protection job.

Reason

The reason the DellEMC Unity storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Replication

The following fields and corresponding data display in the Compliant for Replication section of the DellEMC Unity RPO Compliance report:

Resource

The name of the primary source volume.

Resource Type

The type associated with the DellEMC Unity storage system. Resource types include file system and LUN.

Source

The source of the replication.

Destination

The destination of the replication.

Job Name

The Replication Backup job associated with the Dell EMC Unity storage system.

Protection Time

The most recent instance of a successful run of the Backup job.

Remaining Compliance Time

The time remaining before your Dell EMC Unity storage system will be non-compliant.

RELATED TOPICS:

- [Protection Compliance Reports](#) on page **398**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

File System RPO Compliance Report

The File System RPO Compliance report displays your file systems in relation to your recovery point objective parameters. Determine which of your file systems are not in compliance with your RPO parameters, and discover the reasons for their non-compliance.

BEFORE YOU BEGIN:

- Create and run a File System Inventory job. You can select one or more file system in a single job definition for cataloging. See [Create an Inventory Job Definition - File System](#) on page 173.
- Create and run a File System Backup job. See [Create a Backup Job Definition - File System](#) on page 216.

Parameters

Use the following parameters to customize your report:

- Host
 - Multiple selections are supported.
- Protection Type
 - Set the protection types to return in the report. Values include Primary and Replication. Multiple selects are supported.
- Storage Vendor
 - Set the storage vendor types to display in the report. Multiple selections are supported.
- Display File Systems That Are
 - Set the compliance status of your file system to return in the report. Values include Compliant and Not Compliant. Multiple selections are supported. By default, this parameter is set to Not Compliant.
- RPO Older Than
 - Set the age of the recovery point objective in days.

The default report parameters report on all non-compliant file systems based on an RPO older than one day.

Quick View

The Quick View section displays a bar graph of compliant and non-compliant file systems based on your RPO parameters.

Not Compliant for Primary Protection

The following fields and corresponding data display in the Not Compliant for Primary Protection section of the File System RPO Compliance report:

File System Mount Point(s)

The file system's associated mount points.

Host (OS Type)

The hostname and operating system type. For example, Windows, AIX, or Linux.

Job Name

The Backup job associated with the file system.

Storage Vendor

The storage vendor associated with the file system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Reason

The reason the file system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Primary Protection

The following fields and corresponding data display in the Compliant for Primary Protection section of the File System RPO Compliance report:

File System Mount Point(s)

The file system's associated mount points.

Host (OS Type)

The hostname and operating system type. For example, Windows, AIX, or Linux.

Job Name

The Backup job associated with the file system.

Storage Vendor

The storage vendor associated with the file system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your file system will be non-compliant.

Not Compliant for Replication

The following fields and corresponding data display in the Not Compliant for Replication section of the File System RPO Compliance report:

File System Mount Point(s)

The file system's associated mount points.

Source

The source of the secondary protection.

Destination

The destination of the secondary protection.

Job Name

The Backup job associated with the file system.

Protection Time

The most recent instance of a successful run of the Backup job.

Storage Vendor

The storage vendor associated with the file system.

Reason

The reason the file system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Replication

The following fields and corresponding data display in the Compliant for Replication section of the File System RPO Compliance report:

File System Mount Point(s)

The file system's associated mount points.

Host (OS Type)

The hostname and operating system type. For example, Windows, AIX, or Linux.

Job Name

The Backup job associated with the file system.

Storage Vendor

The storage vendor associated with the file system.

Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your file system will be non-compliant.

RELATED TOPICS:

- [Protection Compliance Reports](#) on page **398**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

HPE Nimble Storage RPO Compliance Report

The HPE Nimble Storage RPO Compliance report displays HPE Nimble Storage systems in relation to your recovery point objective parameters. Determine which of your HPE Nimble Storage systems are not in compliance with your RPO parameters, and discover the reasons for their non-compliance.

BEFORE YOU BEGIN:

- Create and run a HPE Nimble Storage Inventory job. You can select one or more HPE Nimble Storage provider in a single job definition for cataloging. See [Create an Inventory Job Definition - Pure Storage FlashArray](#) on page **192**.
- Create and run a HPE Nimble Storage Backup job. See [Create a Backup Job Definition - Pure Storage FlashArray](#) on page **238**.

Parameters

Use the following parameters to customize your report:

- HPE Nimble Storage
Multiple selections are supported.
- Protection Type
Set the HPE Nimble Storage protection type to return in the report. Values include Primary and Replication. Multiple selections are supported.
- Display Resources That Are
Set the compliance status of your HPE Nimble Storage systems to return in the report. Values include Compliant and Not Compliant. Multiple selections are supported. By default, this parameter is set to Not Compliant.
- RPO Older Than
Set the age of the recovery point objective in days.

Quick View

The Quick View section displays a bar graph of compliant and non-compliant HPE Nimble Storage systems based on your RPO parameters.

Not Compliant for Primary Protection

The following fields and corresponding data display in the Not Compliant for Primary Protection section of the HPE Nimble Storage RPO Compliance report:

Volume

The name of the HPE Nimble Storage system volume.

Location

The location of the HPE Nimble Storage system.

Job Name

The Backup job associated with the HPE Nimble Storage system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Reason

The reason the HPE Nimble Storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Primary Protection

The following fields and corresponding data display in the Compliant for Primary Protection section of the HPE Nimble Storage RPO Compliance report:

Volume

The name of the HPE Nimble Storage system volume.

Location

The location of the HPE Nimble Storage system.

Job Name

The Backup job associated with the HPE Nimble Storage system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your HPE Nimble Storage system will be non-compliant.

Not Compliant for Replication

The following fields and corresponding data display in the Not Compliant for Replication section of the HPE Nimble Storage RPO Compliance report:

Volume

The name of the HPE Nimble Storage system volume.

Source Storage Array

The source of the secondary protection.

Destination Storage Array

The destination of the secondary protection.

Job Name

The Backup job associated with the HPE Nimble Storage storage system.

Last Successful Protection Time

The most recent instance of a successful run of the protection job.

Reason

The reason the HPE Nimble Storage storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Replication

The following fields and corresponding data display in the Compliant for Replication section of the HPE Nimble Storage RPO Compliance report:

Volume

The name of the HPE Nimble Storage system volume.

Source Storage Array

The source of the secondary protection.

Destination Storage Array

The destination of the secondary protection.

Job Name

The Backup job associated with the HPE Nimble Storage storage system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your HPE Nimble Storage storage system will be non-compliant. Hover over the bar to see the remaining compliance time.

RELATED TOPICS:

- [Protection Compliance Reports](#) on page **398**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

IBM Spectrum Accelerate RPO Compliance Report

The IBM Spectrum Accelerate RPO Compliance report displays IBM Spectrum Accelerate storage systems in relation to your recovery point objective parameters. Determine which of your IBM Spectrum Accelerate storage systems are not in compliance with your RPO parameters, and discover the reasons for their non-compliance.

BEFORE YOU BEGIN:

- Create and run an IBM Spectrum Accelerate Inventory job. You can select one or more IBM provider in a single job definition for cataloging. See [Create an Inventory Job Definition - IBM Spectrum Accelerate](#) on page **177**.
- Create and run an IBM Spectrum Accelerate Backup job. See [Create a Backup Job Definition - IBM Spectrum Accelerate](#) on page **223**.

Use the IBM Spectrum Accelerate RPO Compliance report to answer questions such as:

- Which of my IBM storage systems are not RPO compliant for Flash Copy protection?
- Which of my IBM Backup jobs have never run successfully?

Parameters

Use the following parameters to customize your report:

- Storage Array
Multiple selections are supported.
- Display Resources That Are
Set the compliance status of your IBM storage systems to return in the report. Values include Compliant and Not Compliant. Multiple selections are supported. By default, this parameter is set to Not Compliant.
- RPO Older Than
Set the age of the recovery point objective in days.

Quick View

The Quick View section displays a bar graph of compliant and non-compliant IBM storage systems based on your RPO parameters.

Not Compliant for Primary Protection

The following fields and corresponding data display in the Not Compliant for Primary Protection section of the IBM Spectrum Accelerate RPO Compliance report:

Volume

The name of the IBM storage system.

Location

The location of the IBM storage system.

Job Name

The Backup job associated with the IBM storage system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Reason

The reason the IBM storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Primary Protection

The following fields and corresponding data display in the Compliant for Primary Protection section of the IBM Spectrum Accelerate RPO Compliance report:

Volume

The name of the IBM storage system.

Location

The location of the IBM storage system.

Job Name

The Backup job associated with the IBM storage system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your IBM storage system will be non-compliant.

RELATED TOPICS:

- [Protection Compliance Reports](#) on page **398**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

IBM Spectrum Virtualize RPO Compliance Report

The IBM Spectrum Virtualize RPO Compliance report displays IBM storage systems in relation to your recovery point objective parameters. Determine which of your IBM storage systems are not in compliance with your RPO parameters, and discover the reasons for their non-compliance.

BEFORE YOU BEGIN:

- Create and run an IBM Spectrum Virtualize Inventory job. You can select one or more IBM provider in a single job definition for cataloging. See [Create an Inventory Job Definition - IBM Spectrum Virtualize](#) on page **181**.
- Create and run an IBM Spectrum Virtualize Backup job. See [Create a Backup Job Definition - IBM Spectrum Virtualize](#) on page **227**.

Use the IBM Spectrum Virtualize RPO Compliance report to answer questions such as:

- Which of my IBM storage systems are not RPO compliant for Flash Copy protection?
- Which of my IBM Backup jobs have never run successfully?

Parameters

Use the following parameters to customize your report:

- Storage Array
Multiple selections are supported.
- Protection Type
Set the IBM protection type to return in the report. Values include FlashCopy and Global Mirror with Change Volumes. Multiple selections are supported.
- Display Resources That Are
Set the compliance status of your IBM storage systems to return in the report. Values include Compliant and Not Compliant. Multiple selections are supported. By default, this parameter is set to Not Compliant.
- RPO Older Than
Set the age of the recovery point objective in days.

The default report parameters report on all non-compliant IBM storage systems based on an RPO older than one day.

Quick View

The Quick View section displays a bar graph of compliant and non-compliant IBM storage systems based on your RPO parameters.

Not Compliant for Flash Copy Protection

The following fields and corresponding data display in the Not Compliant for Flash Copy Protection section of the IBM Spectrum Virtualize RPO Compliance report:

Volume

The name of the IBM storage system.

Location

The location of the IBM storage system.

Consistency Group

The name of the volume's associated consistency group.

Job Name

The Backup job associated with the IBM storage system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Reason

The reason the IBM storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Flash Copy Protection

The following fields and corresponding data display in the Compliant for Snapshot Protection section of the IBM Spectrum Virtualize RPO Compliance report:

Volume

The name of the IBM storage system.

Location

The location of the IBM storage system.

Consistency Group

The name of the volume's associated consistency group.

Job Name

The Backup job associated with the IBM storage system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your IBM storage system will be non-compliant. Hover over the bar to see the remaining compliance time.

Not Compliant for Global Mirror with Change Volumes Protection

The following fields and corresponding data display in the Not Compliant for Global Mirror with Change Volumes Protection section of the IBM Spectrum Virtualize RPO Compliance report:

Volume

The name of the primary source volume.

Source

The source of the secondary protection.

Consistency Group

The name of the volume's associated consistency group.

Destination

The destination of the secondary protection.

Job Name

The Global Mirror job associated with the IBM storage system.

Protection Time

The most recent instance of a successful run of the protection job.

Reason

The reason the IBM storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Global Mirror with Change Volumes Protection

The following fields and corresponding data display in the Compliant for Global Mirror with Change Volumes Protection section of the IBM Spectrum Virtualize RPO Compliance report:

Volume

The name of the primary source volume.

Source

The source of the secondary protection.

Consistency Group

The name of the volume's associated consistency group.

Destination

The destination of the secondary protection.

Job Name

The Backup job associated with the IBM storage system.

Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your IBM storage system will be non-compliant. Hover over the bar to see the remaining compliance time.

RELATED TOPICS:

- [Protection Compliance Reports](#) on page **398**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

NetApp ONTAP Protection Usage Report

The NetApp ONTAP Protection Space Usage report displays the storage usage of IBM Spectrum Copy Data Management protection jobs. Review the amount of space occupied by snapshots, SnapVaults, and SnapMirrors on the volumes of your NetApp ONTAP storage systems.

BEFORE YOU BEGIN:

- Create and run a NetApp ONTAP Storage Inventory job. You can select one or more NetApp ONTAP cluster providers in a single job definition for cataloging. See [Create an Inventory Job Definition - NetApp ONTAP Storage](#) on page 183.
- Create and run a NetApp ONTAP Backup job. See [Create a Backup Job Definition - NetApp ONTAP](#) on page 231.

Use the NetApp ONTAP Protection Space Usage report to answer questions such as:

- What is the secondary protection storage usage across my volumes?
- What is the combined size of all of my volumes on a NetApp ONTAP storage system?

Parameters

Use the following parameters to customize your report:

- Storage Array
Multiple selections are supported.
- Protection Type
Set the NetApp ONTAP protection type. Protection types include SnapMirror, SnapVault, and Snapshot.
Multiple selections are supported.
- Job Type
Set the IBM Spectrum Copy Data Management protection jobs to display in the report. Multiple selections are supported.

Quick View

The Quick View section displays a bar graph of the storage protection usage across all of the volumes on the selected NetApp ONTAP storage system.

Storage Used by Snapshot Protection

The following fields and corresponding data display in the Storage Used by Snapshot Protection section of the NetApp ONTAP Protection Usage report:

Primary Source Volume

The name of the primary source volume.

Location

The location of the primary source volume.

Job Names

The names of the associated IBM Spectrum Copy Data Management protection jobs.

Snapshot Count

The number of snapshots available on the volume.

Oldest Snapshot Creation Time

The creation date and time of the oldest snapshot on the volume

Volume Size

The total size of the volume.

Volume Used Size

The size of the volume occupied by data.

Snapshot Usage

The amount of space on the volume dedicated to snapshot protection.

Storage Used by SnapVault/SnapMirror Protection

The following fields and corresponding data display in the Storage Used by SnapVault/SnapMirror Protection section of the NetApp ONTAP Protection Usage report:

Primary Source Volume

The name of the primary source volume.

Destination

The destination of the secondary protection.

Job Names

The names of the associated IBM Spectrum Copy Data Management protection jobs.

SnapVault Usage

The amount of space on the volume dedicated to SnapVault protection.

SnapMirror Usage

The amount of space on the volume dedicated to SnapMirror protection.

RELATED TOPICS:

- [Protection Compliance Reports](#) on page 398
- [Report Overview](#) on page 360

- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

NetApp ONTAP RPO Compliance Report

The NetApp ONTAP RPO Compliance report displays NetApp ONTAP storage systems in relation to your recovery point objective parameters. Determine which of your NetApp ONTAP storage systems are not in compliance with your RPO parameters, and discover the reasons for their non-compliance.

BEFORE YOU BEGIN:

- Create and run a NetApp ONTAP Storage Inventory job. You can select one or more NetApp ONTAP cluster providers in a single job definition for cataloging. See [Create an Inventory Job Definition - NetApp ONTAP Storage](#) on page 183.
- Create and run a NetApp ONTAP Backup job. See [Create a Backup Job Definition - NetApp ONTAP](#) on page 231.

Use the NetApp ONTAP RPO Compliance report to answer questions such as:

- Which of my NetApp ONTAP storage systems are not RPO compliant for SnapVault or SnapMirror protection?
- Which of my NetApp ONTAP Backup jobs have never run successfully?

Parameters

Use the following parameters to customize your report:

- Storage Array
 - Multiple selections are supported.
- Protection Type
 - Set the RPO compliance protection type. Protection types include SnapMirror, SnapVault, and Snapshot. Multiple selections are supported. By default, this parameter is set to All.
- Display Resources That Are
 - Set the compliance status of your NetApp ONTAP storage systems to return in the report. Values include Compliant and Not Compliant. Multiple selections are supported. By default, this parameter is set to Not Compliant.
- RPO Older Than
 - Set the age of the recovery point objective in days.

The default report parameters report on all non-compliant NetApp ONTAP storage systems based on an RPO older than one day.

Quick View

The Quick View section displays a bar graph of compliant and non-compliant NetApp ONTAP storage systems based on your RPO parameters.

Not Compliant for Primary Protection

The following fields and corresponding data display in the Not Compliant for Primary Protection section of the NetApp ONTAP RPO Compliance report:

Resource Name

The name of the NetApp ONTAP storage system.

Location

The location of the NetApp ONTAP storage system.

Job Name

The Backup job associated with the NetApp ONTAP storage system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Reason

The reason the NetApp ONTAP storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Primary Protection

The following fields and corresponding data display in the Compliant for Primary Protection section of the NetApp ONTAP RPO Compliance report:

Resource Name

The name of the NetApp ONTAP storage system.

Location

The location of the NetApp ONTAP storage system.

Job Name

The Backup job associated with the NetApp ONTAP storage system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your NetApp ONTAP storage system will be non-compliant. Hover over the bar to see the remaining compliance time.

Not Compliant for Replication

The following fields and corresponding data display in the Not Compliant for Replication section of the NetApp ONTAP RPO Compliance report:

Source

The source of the secondary protection.

Destination

The destination of the secondary protection.

Job Name

The SnapVault or SnapMirror Backup job associated with the NetApp ONTAP storage system.

Protection Time

The most recent instance of a successful run of the protection job.

Reason

The reason the NetApp ONTAP storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Replication

The following fields and corresponding data display in the Compliant for Replication section of the NetApp ONTAP RPO Compliance report:

Source

The source of the secondary protection.

Destination

The destination of the secondary protection.

Job Name

The Backup job associated with the NetApp ONTAP storage system.

Protection Time

The most recent instance of a successful run of the Backup job.

Remaining Compliance Time

The time remaining before your NetApp ONTAP storage system will be non-compliant. Hover over the bar to see the remaining compliance time.

RELATED TOPICS:

- [Protection Compliance Reports](#) on page **398**
- [Report Overview](#) on page **360**

- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Pure Storage FlashArray RPO Compliance Report

The Pure Storage FlashArray RPO Compliance report displays Pure Storage systems in relation to your recovery point objective parameters. Determine which of your Pure Storage systems are not in compliance with your RPO parameters, and discover the reasons for their non-compliance.

BEFORE YOU BEGIN:

- Create and run a Pure Storage Inventory job. You can select one or more Pure Storage provider in a single job definition for cataloging. See [Create an Inventory Job Definition - Pure Storage FlashArray](#) on page 192.
- Create and run a Pure Storage Backup job. See [Create a Backup Job Definition - Pure Storage FlashArray](#) on page 238.

Parameters

Use the following parameters to customize your report:

- Storage Array
Multiple selections are supported.
- Protection Type
Set the Pure Storage FlashArray protection type to return in the report. Values include Primary and Replication. Multiple selections are supported.
- Display Resources That Are
Set the compliance status of your Pure Storage systems to return in the report. Values include Compliant and Not Compliant. Multiple selections are supported. By default, this parameter is set to Not Compliant.
- RPO Older Than
Set the age of the recovery point objective in days.

Quick View

The Quick View section displays a bar graph of compliant and non-compliant Pure Storage FlashArray systems based on your RPO parameters.

Not Compliant for Primary Protection

The following fields and corresponding data display in the Not Compliant for Primary Protection section of the Pure Storage FlashArray RPO Compliance report:

Volume

The name of the Pure Storage FlashArray system volume.

Location

The location of the Pure Storage FlashArray system.

Job Name

The Backup job associated with the Pure Storage FlashArray system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Reason

The reason the Pure Storage FlashArray system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Primary Protection

The following fields and corresponding data display in the Compliant for Primary Protection section of the Pure Storage FlashArray RPO Compliance report:

Volume

The name of the Pure Storage FlashArray system volume.

Location

The location of the Pure Storage FlashArray system.

Job Name

The Backup job associated with the Pure Storage FlashArray system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your Pure Storage FlashArray system will be non-compliant.

Not Compliant for Replication

The following fields and corresponding data display in the Not Compliant for Replication section of the Pure Storage FlashArray RPO Compliance report:

Volume

The name of the Pure Storage FlashArray system volume.

Source Storage Array

The source of the secondary protection.

Destination Storage Array

The destination of the secondary protection.

Job Name

The Backup job associated with the Pure Storage FlashArray storage system.

Last Successful Protection Time

The most recent instance of a successful run of the protection job.

Reason

The reason the Pure Storage FlashArray storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Replication

The following fields and corresponding data display in the Compliant for Replication section of the Pure Storage FlashArray RPO Compliance report:

Volume

The name of the Pure Storage FlashArray system volume.

Source Storage Array

The source of the secondary protection.

Destination Storage Array

The destination of the secondary protection.

Job Name

The Backup job associated with the Pure Storage FlashArray storage system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your Pure Storage FlashArray storage system will be non-compliant. Hover over the bar to see the remaining compliance time.

RELATED TOPICS:

- [Protection Compliance Reports](#) on page **398**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Recovery Points Report

The Recovery Points report displays recovery points of protected resources. View the protection type, job name, and location of your available recovery points.

BEFORE YOU BEGIN:

- Create and run a Backup job.

Parameters

Use the following parameters to customize your report:

- Site
 - Multiple selections are supported.
- Resource Type
 - Set the resource type to return in the report. Resource types include Oracle Database, SAP HANA Database, SQL Database, VMware Datastores, and VMware VMs. By default, this parameter is set to All.
- Protection Type
 - Set the protection type to display in the report. Protection types include primary and secondary. Multiple selections are supported. By default, this parameter is set to All.
- Show Recovery Points
 - Set the age of the recovery point objective in days.

Recovery Points

The following fields and corresponding data display in the Recovery Points section of the Recovery Points report:

Resource Name

The name and location of the protected resource.

Site

The site associated with the protected resource.

Protection Time

The most recent instance of a successful run of the Backup job.

Protection Type

The recovery point's protection type.

Job Name

The Backup job associated with the recovery point.

Volumes (Location)

The volumes associated with the displayed recovery point.

RELATED TOPICS:

- [Protection Compliance Reports](#) on page **398**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Unprotected Virtual Machines Report

The Unprotected Virtual Machines report displays virtual machines that are not protected with a primary snapshot or VM Replication. Quickly view all of your unprotected virtual machines, along with their eligibility for protection.

BEFORE YOU BEGIN:

- Create and run a VMware Inventory job. See [Create an Inventory Job Definition - VMware](#) on page **194**.
- All related storage providers must be added to IBM Spectrum Copy Data Management, which include NetApp storage controllers and clusters. Create and run a NetApp ONTAP Storage Inventory job to catalog associated NetApp ONTAP storage providers. See [Create an Inventory Job Definition - NetApp ONTAP Storage](#) on page **183**.
- Create and run a VMware Backup job. See [Create a Backup Job Definition - VMware](#) on page **242**.

Use the Unprotected Virtual Machines report to answer questions such as:

- How many of my virtual machines are eligible for primary Storage Snapshot protection?
- How much storage space is used within a given time period?

Parameters

Use the following parameters to customize your report:

- vCenter
Multiple selections are supported.
- Power State
Set the power state of virtual machines returned by the report. Power states include Powered On, Powered Off, Suspended, or All. Multiple selections are supported. By default, this parameter is set to All.

The default report parameters report on unprotected virtual machines on all vCenters, in any power state.

Unprotected Virtual Machines

The following fields and corresponding data display in the Unprotected Virtual Machines section of the Unprotected Virtual Machines report for eligible virtual machines:

VM Name

The name of the virtual machine, preceded by an icon that displays the virtual machine's power state.

Location

The location of the virtual machine.

Hostname

The host node where the virtual machine resides.

Operating System

The operating system associated with the virtual machine

Provisioned Space

The amount of space on the datastore allocated for virtual disk files.

Datastores

The name of the associated datastore.

RELATED TOPICS:

- [Protection Compliance Reports](#) on page **398**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

VMware RPO Compliance Report

The VMware RPO Compliance report displays VMware objects in relation to your recovery point objective parameters. Determine which of your VMware objects are not in compliance with your RPO parameters, and discover the reasons for their non-compliance.

BEFORE YOU BEGIN:

- Create and run a VMware Inventory job. See [Create an Inventory Job Definition - VMware](#) on page **194**.
- All related NetApp ONTAP storage providers must be added to IBM Spectrum Copy Data Management, which include NetApp ONTAP storage controllers and clusters. Create and run a NetApp ONTAP Storage Inventory job to catalog associated NetApp ONTAP storage providers. See [Create an Inventory Job Definition - NetApp ONTAP Storage](#) on page **183**.
- Create and run a VMware Backup job. See [Create a Backup Job Definition - VMware](#) on page **242**.

Use the VMware RPO Compliance report to answer questions such as:

- Which of my VMware Backup job have never run successfully?
- What is the remaining compliance time of a specific VMware object?

Parameters

Use the following parameters to customize your report:

- vCenter
Multiple selections are supported.
- Primary Protection Resource
Set the primary protection resource type to display in the report. Resource types include Datastore and Virtual Machine. Multiple selections are supported. By default, this parameter is set to All.
- Replication Resource
Set the replication resource type to display in the report. Resource types include Datastore and Virtual Machine. Multiple selections are supported. By default, this parameter is set to All.
- Protection Type
Set the RPO compliance protection type. Protection types include Primary and Replication. Multiple selections are supported. By default, this parameter is set to All.
- Storage Vendor
Set the storage vendor types to display in the report. Multiple selections are supported.

- **Display Resources That Are**
Set the compliance status of your VMware objects to return in the report. Values include Compliant and Not Compliant. Multiple selections are supported. By default, this parameter is set to Not Compliant.
- **RPO Older Than**
Set the age of the recovery point objective in days.
- **View VMs with Application(s) only**
Select Yes to view only application virtual machines. Select the application type through the Application Server parameter. Select No to display all virtual machines.

The default report parameters report on all non-compliant VMware objects based on an RPO older than one day.

Quick View

The Quick View section displays a bar graph of compliant and non-compliant VMware objects based on your RPO parameters.

Not Compliant for Primary Protection

The following fields and corresponding data display in the Not Compliant for Primary Protection section of the VMware RPO Compliance report:

Resource Name

The name of the virtual machine, preceded by an icon that displays the object type or a virtual machine's power state.

Location

The location of the VMware object.

Job Name

The Backup job associated with the VMware object.

Storage Vendor

The storage vendor associated with the VMware object.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Application(s)

The application server type, if applicable. Available types include Active Directory Server, Exchange Server, SQL Server, and SharePoint Server.

Reason

The reason the VMware object is not compliant for primary protection. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Primary Protection

The following fields and corresponding data display in the Compliant for Primary Protection section of the VMware RPO Compliance report:

Resource Name

The name of the object, preceded by an icon that displays the object type or a virtual machine's power state.

Location

The location of the VMware object.

Job Name

The Backup job associated with the VMware object.

Storage Vendor

The storage vendor associated with the VMware object.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Application(s)

The application server type, if applicable. Available types include Active Directory Server, Exchange Server, SQL Server, and SharePoint Server.

Compliance Time Remaining

The time remaining before your VMware object will be non-compliant.

Not Compliant for Replication

The following fields and corresponding data display in the Not Compliant for Replication section of the VMware RPO Compliance report:

Resource Name

The name of the virtual machine.

Source

The source of the secondary protection.

Destination

The destination of the secondary protection.

Job Name

The Backup job associated with the VMware object.

Protection Time

The most recent instance of a successful run of the Backup job.

Storage Vendor

The storage vendor associated with the VMware object.

Reason

The reason the VMware object is not compliant for secondary protection. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Replication

The following fields and corresponding data display in the Compliant for Replication section of the VMware RPO Compliance report:

Resource Name

The name of the virtual machine.

Source

The source of the secondary protection.

Destination

The destination of the secondary protection.

Job Name

The Backup job associated with the VMware object.

Protection Time

The most recent instance of a successful run of the Backup job.

Storage Vendor

The storage vendor associated with the VMware object.

Compliance Time Remaining

The time remaining before your VMware object will be non-compliant.

RELATED TOPICS:

- [Protection Compliance Reports](#) on page **398**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Storage Protection Reports

The Storage Protection Reports help ensure your data is protected and display the status of your replication process. Use the Storage Protection Reports to view qtrees that are protected using NetApp ONTAP SnapVault or NetApp ONTAP SnapMirror functionality and volumes that are protected using NetApp ONTAP SnapMirror functionality. You can also view unprotected qtrees and volumes to help plan your data protection strategies.

Reports are based on the data collected by the most recently run job. You must catalog all volumes on a storage system using full storage system cataloging to view the correct storage protection status for all volumes and qtrees.

Use the Storage Protection Reports to answer questions such as:

- What are my protected qtrees and volumes?
- Are there unprotected qtrees and volumes?
- What is the status of my replication process?

Choose the Storage Protection report that fits your needs.

- **NetApp ONTAP OSSV Relationship Status Report** - Review the relationship status of your OSSV protected clients.
- **NetApp ONTAP Overprotected Volumes** - View volumes that are overprotected using NetApp ONTAP SnapVault or SnapMirror functionality.
- **NetApp ONTAP Qtree Protection Status Report** - View qtrees that are protected using NetApp ONTAP SnapVault or SnapMirror functionality.
- **NetApp ONTAP Underprotected Volumes** - View volumes that are underprotected using NetApp ONTAP SnapVault, SnapMirror, or Snapshot functionality.
- **NetApp ONTAP Volume Protection Status Report** - View volumes that are protected using NetApp ONTAP SnapVault, SnapMirror, or Snapshot functionality and evaluate your volume replication processes.
- **NetApp ONTAP Transition Dependency** - View inter-node replication dependencies between sources and destinations for 7-mode systems.

WHY IT MATTERS:

IBM Spectrum Copy Data Management provides reports that identify volumes that are not being protected or where protection may have failed due to an errant configuration change or other issue. With this information, you can investigate the situation and apply proper corrective measures.

RELATED TOPICS:

- [NetApp ONTAP OSSV Relationship Status Report](#) on page 437
- [NetApp ONTAP Overprotected Volumes Report](#) on page 439
- [NetApp ONTAP Qtree Protection Status Report](#) on page 441

- [NetApp ONTAP Underprotected Volumes Report](#) on page **443**
- [NetApp ONTAP Volume Protection Status Report](#) on page **445**
- [NetApp ONTAP Transition Dependency Report](#) on page **447**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**
- [Download a Report](#) on page **365**

NetApp ONTAP OSSV Relationship Status Report

The NetApp ONTAP OSSV Relationship Status report displays the status of OSSV protected clients. An OSSV backup and recovery solution transfers data from an OSSV host to a NetApp ONTAP secondary storage system as a block-level incremental backup.

Use the NetApp ONTAP OSSV Relationship Status report to answer questions such as:

- Which of my OSSV protected clients have not been backed up in more than 30 days?
- What is the backup destination of an OSSV source?

Parameters

Use the following parameters to customize your report:

- OSSV Primary Node
Multiple selections are supported.
- Days Since Last Backup
- Record Limit

The default report parameters display the relationship status of all OSSV nodes.

Detail View - OSSV Protected Clients

The following fields and corresponding data display in the Detail View - OSSV Protected Clients section of the NetApp ONTAP OSSV Relationship Status report:

Source

The node and path where the replication source is located.

Destination Qtree

The node and path of the destination qtree.

Status

The status of the OSSV backup (for example, idle and quiescing).

Latest Backup

The date and time of the most recent OSSV backup.

RELATED TOPICS:

- [Storage Protection Reports](#) on page 435
- [Report Overview](#) on page 360

- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

NetApp ONTAP Overprotected Volumes Report

The NetApp ONTAP Overprotected Volumes report displays volumes that are overprotected using NetApp ONTAP Snapshot, SnapVault, or SnapMirror software. Through your parameters selections, you can decide what constitutes an overprotected volume. Select the number of acceptable snapshots to reside on a volume, and if a volume should also be SnapMirrored and SnapVaulted. If a volume exceeds any of your parameter selections, it will be returned on the NetApp ONTAP Overprotected Volumes report.

Note: This report may discover storage systems that have not been cataloged in the database. These non-cataloged storage systems are discovered due to their replication relationships with other cataloged storage systems.

Use the NetApp ONTAP Overprotected Volumes report to answer questions such as:

- How many of my volumes are overprotected?
- How many gigabytes of space are dedicated to overprotecting a volume?
- Which of my volumes are protected with more than 10 snapshots?

Parameters

Use the following parameters to customize your report:

- Storage Array
 - Multiple selections are supported.
- Acceptable Snapshots
- Is Volume SnapMirrored?
- Is Volume SnapVaulted?

The default report parameters display all overprotected volumes with ten or more snapshots.

Detail View - Overprotected Volumes

The following fields and corresponding data display in the Detail View - Overprotected Volumes section of the NetApp ONTAP Overprotected Volumes report:

Volume

The name of the overprotected volume.

Storage Array

The host where the overprotected volume is located.

Location

The node where the volume resides.

Overprotection Reason

The reason the volume was returned by the report. For example, if the number of snapshots is larger than the defined Acceptable Snapshots parameter, or if a volume is SnapVaulted and the Is Volume SnapVaulted parameter is set to No.

Overprotection Storage Cost

The amount of space on the volume dedicated to overprotection.

RELATED TOPICS:

- [Storage Protection Reports](#) on page **435**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

NetApp ONTAP Qtree Protection Status Report

The NetApp ONTAP Qtree Protection Status report displays qtrees that are protected using NetApp ONTAP Snapshot, SnapVault, or SnapMirror software.

Note: This report may discover storage systems that have not been cataloged in the database. These non-cataloged storage systems are discovered due to their replication relationships with other cataloged storage systems.

Use the NetApp ONTAP Qtree Protection Status report to answer questions such as:

- Which of my cataloged SnapVaulted or SnapMirrored Qtrees are in an unprotected state?
- Which of my NetApp ONTAP volumes are exceeding their lag time by 30 days or more?

Parameters

Use the following parameters to customize your report:

- Storage Array
 - Multiple selections are supported.
- Protection Type
- Lag time in days
- Show Protected Qtrees
- Record Limit

The default report parameters display the unprotected qtrees for all storage systems and all volumes.

Detail View

The following fields and corresponding data display in the following sections: NetApp ONTAP SnapMirrored Qtrees in an Unprotected State, Qtrees Protected by SnapMirror, NetApp ONTAP SnapVaulted Qtrees in an Unprotected State, and Qtrees Protected by SnapVault.

Storage Array

The physical server where your files are stored.

Source

The name of the protected or unprotected qtree.

Destination

The node where the replication destination is located.

State

The state of the destination (for example, SnapMirrored, SnapVaulted, broken-off, uninitialized, or unknown).

Lag Time

The lag time in days, hours, minutes, and seconds between the source and the destination.

Xfer Throughput

The transfer throughput in KBs per second between the source and the destination.

RELATED TOPICS:

- [Storage Protection Reports](#) on page **435**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

NetApp ONTAP Underprotected Volumes Report

The NetApp ONTAP Underprotected Volumes report displays volumes that are underprotected using NetApp ONTAP Snapshot, SnapVault, or SnapMirror software. Through your parameters selections, you can decide what constitutes an underprotected volume. Select the number of acceptable snapshots to reside on a volume, the lag time, and if a volume should also be SnapMirrored and SnapVaulted. If a volume has less protection than your parameter selections, it will be returned on the NetApp ONTAP Underprotected Volumes report.

Note: This report may discover storage systems that have not been cataloged in the database. These non-cataloged storage systems are discovered due to their replication relationships with other cataloged storage systems.

Use the NetApp ONTAP Underprotected Volumes report to answer questions such as:

- What are the names of my underprotected volumes?
- Which of my volumes are underprotected with less than 10 snapshots?

Parameters

Use the following parameters to customize your report:

- Storage Array
 - Multiple selections are supported.
- Acceptable Snapshots
- Lag time in days
- Is Volume SnapMirrored?
- Is Volume SnapVaulted?

The default report parameters display underprotected volumes with less than ten snapshots that are older than 7 days.

Underprotected Volumes

The following fields and corresponding data display in the Underprotected Volumes section of the NetApp ONTAP Underprotected Volumes report:

Volume

The name of the underprotected volume.

Storage Array

The node where the underprotected volume is located.

Location

The node where the volume resides.

Underprotection Reason

The reason the volume was returned by the report. For example, if the number of snapshots is lower than the defined No. of Acceptable Snapshots parameter, or if the snapshots have excessive lag times.

RELATED TOPICS:

- [Storage Protection Reports](#) on page **435**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

NetApp ONTAP Volume Protection Status Report

The NetApp ONTAP Volume Protection Status report displays volumes that are protected or unprotected using NetApp ONTAP Snapshot, SnapVault, or SnapMirror software.

Note: This report may discover storage systems that have not been cataloged in the database. These non-cataloged storage systems are discovered due to their replication relationships with other cataloged storage systems.

Use the NetApp ONTAP Volume Protection Status report to answer questions such as:

- Which of my cataloged SnapVaulted or SnapMirrored volumes are in an unprotected state?
- Which of my NetApp ONTAP volumes are exceeding their lag time by 30 days or more?

Parameters

Use the following parameters to customize your report:

- Storage Array
 - Multiple selections are supported.
- Protection Type
- Lag time in days
- Show Protected Volumes

The default report parameters display all unprotected SnapMirror volumes.

Detail View - NetApp ONTAP SnapMirrored and SnapVaulted Volumes

The following fields and corresponding data display in the following sections: NetApp ONTAP SnapMirrored Volumes in an Unprotected State, NetApp ONTAP Volumes Protected by SnapMirror, NetApp ONTAP SnapVaulted Volumes in an Unprotected State and NetApp ONTAP Volumes Protected by SnapVault.

Storage Array

The node on which the protected or unprotected volume is located.

Source

The name of the protected or unprotected volume.

Destination

The node on which the replication destination is located.

State

The state of the destination. For example, SnapMirrored, SnapVaulted, broken-off, uninitialized, or unknown.

Type

The SnapMirror type for clustered volumes. For example, Vault or Mirror. Blank entries indicate non-clustered 7-Mode volumes.

Lag Time

The lag time in days, hours, minutes, and seconds between the source and the destination.

Xfer Throughput (NetApp ONTAP Volumes Protected by SnapMirror/SnapVault section only)

The transfer throughput in KBs per second between the source and the destination.

Detail View - NetApp ONTAP Snapshotted Volumes

The following fields and corresponding data display in the following sections NetApp ONTAP Volume with Snapshots Exceeding Lag Time and NetApp ONTAP Volumes Protected by Snapshot.

Storage Array

The node on which the protected or unprotected volume is located.

Volume

The name of the protected or unprotected volume.

Location

The node where the volume resides.

Snapshot Count

The number of snapshots available on the volume.

Latest Snapshot Time

The date and time of the most recent snapshot on the volume.

RELATED TOPICS:

- [Storage Protection Reports](#) on page **435**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

NetApp ONTAP Transition Dependency Report

The NetApp ONTAP Transition Dependency report displays the destination nodes and replication relationships for all of your NetApp ONTAP storage systems operating in 7-mode.

BEFORE YOU BEGIN:

- Create and run a NetApp ONTAP Storage Inventory job to catalog your NetApp ONTAP storage systems operating in 7-mode. You can select one or more NetApp ONTAP cluster provider in a single job definition for cataloging. See [Create an Inventory Job Definition - NetApp ONTAP Storage](#) on page 183.
- Create and run a NetApp ONTAP Backup job. See [Create a Backup Job Definition - NetApp ONTAP](#) on page 231.

Parameters

Use the following parameters to customize your report:

- Storage Array
Set the NetApp ONTAP storage systems to display in the report. Multiple selections are supported.-
- Show Detailed View
Enable to show detailed volume, SnapMirror, and SnapVault data in the report.

Summary View

The following fields and corresponding data display in the Summary View section of the NetApp ONTAP Transition Dependency report:

Storage Array

The name of the NetApp ONTAP storage system source, along with the operating system version running on the source node.

SnapMirror Destination Nodes (Count)

The name of the SnapMirror destination node. The relationship count displays in parentheses.

SnapVault Destination Nodes (Count)

The name of the SnapVault destination node. The relationship count displays in parentheses.

Destination Node OS Version

The operating system version running on the destination node.

Detail View

If the Show Detailed View parameter is set to Yes, the following fields and corresponding data display in the Detail View section of the NetApp ONTAP Transition Dependency report:

Storage Array: Volume

The name of the NetApp ONTAP storage system and associated source volume, along with the size of the source volume.

SnapMirror Destination

The name of the SnapMirror destination and associated volume.

SnapVault Destination

The name of the SnapVault destination and associated volume.

RELATED TOPICS:

- [Storage Protection Reports](#) on page **435**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Storage Utilization Reports

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

Use the Storage Utilization Reports to answer questions such as:

- What is the total available storage space in the entire system?
- What is the amount of free and used space on my volumes?
- How many files and disks make up my aggregates?

The information in these reports are presented in a chart-based **Quick View** section, or tabular **Summary View** and **Detail View** sections.

Choose the Storage Utilization report that fits your needs.

- **DellEMC Unity File Systems Report** - Review the storage utilization of your DellEMC Unity file systems.
- **DellEMC Unity LUNs Report** - Review the total capacity of your LUNs, the total free space, and the percentage available to ascertain your DellEMC Unity LUN storage utilization.
- **DellEMC Unity Pools Report** - Review the storage utilization of your DellEMC Unity pools. View the total and free space available and the number of volumes and disks that make up your storage pools.
- **HPE Nimble Storage Volumes Report** - Review the total capacity of your volumes, the total free space, and the percentage available to ascertain your HPE Nimble Storage volume storage utilization.
- **IBM Spectrum Accelerate Pools Report** - Review the storage utilization of your IBM Spectrum Accelerate Pools. View the total and free space available and the number of volumes and disks that make up your pools.
- **IBM Spectrum Accelerate Volumes Report** - View the total space used by and free space available to your IBM Spectrum Accelerate volumes.
- **IBM Spectrum Virtualize Consistency Groups Report** - Display information about your IBM Consistency Groups. View the associated storage providers, source and target volumes, and protection status of your IBM volumes through Consistency Groups
- **IBM Spectrum Virtualize Pools Report** - Review the storage utilization of your IBM storage pools. View the total and free space available and the number of volumes and disks that make up your storage pools.
- **IBM Spectrum Virtualize Volumes Report** - Review the storage utilization of your IBM volumes. View the total space consumed as well as the free space available on your IBM volumes.

WHY IT MATTERS:

Utilization reports provide advanced warning for volumes, aggregates, or LUNs that are beyond a specified capacity range. Similar reports are available for VMware datastores.

- **Instant Disk Restore Volumes Report** - Display a list of mapped volumes created through Instant Disk Restore jobs.
- **NetApp ONTAP Aggregates Report** - Review the storage utilization and configuration of your NetApp aggregates. View the total and free space available and the number of volumes and disks that make up your aggregates.
- **NetApp ONTAP LUNs Report** - Review the total capacity of your LUNs, the total free space, and the percentage available to ascertain your NetApp ONTAP LUN storage utilization.
- **NetApp ONTAP Orphaned LUNs Report** - Review NetApp ONTAP storage orphaned LUNs. These are the LUNs that have no initiator group mapping or belong to volumes that are offline.
- **NetApp ONTAP Quotas Report** - Review quota status to determine which users or groups are approaching or exceeding quota limits.
- **NetApp ONTAP Snapshots Report** - Review the total capacity of your snapshots, the total free space, and the percentage available to ascertain your NetApp ONTAP Snapshot storage utilization.
- **NetApp ONTAP Volumes Report** - Review the total capacity of your volumes, the total free space, and the percentage available to ascertain your NetApp ONTAP volume storage utilization.
- **Pure Storage FlashArray Volumes Report** - Review the total capacity of your volumes, the total free space, and the percentage available to ascertain your Pure Storage FlashArray volume storage utilization.
- **Storage Capacity Report** - Report the storage capacity of your IBM Spectrum Virtualize pools, DellEMC Unity pools, and NetApp ONTAP aggregates.
- **VM and Storage Mapping Report** - Report that maps VMs all the way to the physical storage from which the datastore is created.
- **VMware Datastores Report** - Review the total capacity of your datastores, the total free space, and the percentage available to ascertain your VMware datastore storage utilization.
- **VMware LUNs Report** - Displays information about VMware LUNs such as which ESX server it belongs to, its datastore, vendor, total, and allocated capacity.
- **VMware Orphaned Datastores Report** - Review the datastores that do not have any virtual machines assigned to them, or if virtual machines are assigned to the datastores, view the virtual machines that are in an inaccessible state.
- **VMware Orphaned LUNs Report** - Review VMware orphaned LUNs. These are the LUNs not used as datastores or RDMs.
- **VMware VM Snapshot Sprawl Report** - Displays information about virtual machines with aged and memory snapshots.
- **VMware VM Sprawl Report** - Displays storage utilization across virtual machines based on their power state and storage utilization across virtual machine templates.
- **VMware VM Storage Report** - Review your virtual machines and associated datastores.

Quick View

This area of the report is a graphical illustration of the report using pie charts. For example, the quick view of the NetApp ONTAP Storage Volumes report shows the total capacity of your volumes, the free space, and the used space.

Summary View

This area of the report displays a summary of the data returned in the report. For example, the summary view of the NetApp ONTAP Storage Aggregates report shows the total used, free, and reserved space on your aggregate.

Detail View

This area of the report is a table where each row details a storage system, its corresponding volume or aggregate, and details returned by the report. For example, the NetApp ONTAP Storage Aggregates report shows the used and free space, volume count, disk count, and status of your aggregates.

Choose the Storage Utilization Report that fits your needs.

RELATED TOPICS:

- [DellEMC Unity File Systems Report](#) on page **453**
- [DellEMC Unity LUNs Report](#) on page **455**
- [DellEMC Unity Pools Report](#) on page **458**
- [HPE Nimble Storage Volumes Report](#) on page **460**
- [IBM Spectrum Accelerate Pools Report](#) on page **463**
- [IBM Spectrum Accelerate Volumes Report](#) on page **465**
- [IBM Spectrum Virtualize Consistency Groups Report](#) on page **467**
- [IBM Spectrum Virtualize Pools Report](#) on page **469**
- [IBM Spectrum Virtualize Volumes Report](#) on page **472**
- [Instant Disk Restore Volumes Report](#) on page **475**
- [NetApp ONTAP Aggregates Report](#) on page **476**
- [NetApp ONTAP LUNs Report](#) on page **479**
- [NetApp ONTAP Orphaned LUNs Report](#) on page **482**
- [NetApp ONTAP Quotas Report](#) on page **484**
- [NetApp ONTAP Snapshots Report](#) on page **486**
- [NetApp ONTAP Volumes Report](#) on page **488**
- [Pure Storage FlashArray Volumes Report](#) on page **491**
- [Storage Capacity Report](#) on page **494**
- [VM and Storage Mapping Report](#) on page **496**

- [VMware Datastores Report](#) on page **498**
- [VMware LUNs Report](#) on page **501**
- [VMware Orphaned Datastores Report](#) on page **503**
- [VMware Orphaned LUNs Report](#) on page **505**
- [VMware VM Snapshot Sprawl Report](#) on page **507**
- [VMware VM Sprawl Report](#) on page **509**
- [VMware VM Storage Report](#) on page **512**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**
- [Download a Report](#) on page **365**

DellEMC Unity File Systems Report

Run the DellEMC Unity File Systems report to display the storage utilization of your DellEMC Unity storage pools. View the number of file systems on your DellEMC Unity storage system, the amount of storage reserved for primary data, and the number of snapshots for the file system.

BEFORE YOU BEGIN:

- Create and run a DellEMC Unity Inventory job. You can select one or more DellEMC Unity providers in a single job definition for cataloging. See [Create an Inventory Job Definition - DellEMC Unity](#) on page 175.

Parameters

Use the following parameters to customize your report:

- Storage Array
 - Multiple selections are supported.
- Detail View Filter
 - Select the alert threshold percentage to display in the report. Select Any to view the details of every DellEMC Unity file system.

Summary View

The following fields and corresponding data display in the Summary View section of the DellEMC Unity File Systems report:

DellEMC Unity Storage

The physical server where your files are stored.

File Systems Count

The number of file systems on the storage system.

Total Size

The quantity of storage reserved for primary data.

Used

The quantity of primary storage allocated.

Free

The quantity of primary storage that is unallocated.

Detail View

The following fields and corresponding data display in the Detail View section of the DellEMC Unity File Systems report:

File System

The name of the file system.

DellEMC Unity Storage

The physical server where your files are stored.

Type

The file system type.

Storage Pool

The name of the associated storage pool.

Total Size

The quantity of storage reserved for primary data.

Used

The quantity of primary storage allocated.

Free

The quantity of primary storage that is unallocated.

No. of Snapshots

The number of snapshots for the file system.

Status

The current status of the file system.

Thin

The disk format of the file system, either thick or thin provisioned.

% Used/Free

A status bar that displays the used space on the file system.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

DellEMC Unity LUNs Report

Review the storage utilization of your DellEMC Unity LUNs. Run the DellEMC Unity LUNs report to view your LUNs, total and available space, and online status.

BEFORE YOU BEGIN:

- Create and run a DellEMC Unity Inventory job. You can select one or more DellEMC Unity providers in a single job definition for cataloging. See [Create an Inventory Job Definition - DellEMC Unity](#) on page 175.

Use the DellEMC Unity LUNs report to answer questions such as:

- How many LUNs are associated with a storage system?
- What is the total and allocated capacity of a LUN?
- What LUNs are not being used, so I can reclaim this space?

Parameters

Use the following parameters to customize your report:

- Storage Array
- Detail View Filter

The default report parameters report on all LUNs with more than 80% space used.

Quick View

The Quick View section displays a pie chart of used and free space on your LUNs. Use the DellEMC Unity parameter to display LUNs on all storage systems or a specific storage system.

Note: The Quick View section is only modified through the DellEMC Unity Storage parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the DellEMC Unity LUNs report:

DellEMC Unity Storage

The physical server where your files are stored.

LUNs Count

The number of LUNs on the storage system.

Total Size

The total size of the LUNs on the storage system.

Allocated

The quantity of primary storage allocated.

Unallocated

The quantity of primary storage that is unallocated.

Detail View

The following fields and corresponding data display in the Detail View section of the DellEMC Unity LUNs report:

LUN

The name of the LUN.

DellEMC Unity Storage

The physical server where your files are stored.

Type

The LUN type. Types include Standalone, VmWareISCSI, and Generic Storage.

Storage Pool

The name of the associated storage pool.

Total Size

The quantity of storage reserved for primary data.

Allocated

The quantity of primary storage allocated.

Unallocated

The quantity of primary storage that is unallocated.

No. of Snapshots

The number of snapshots for the LUN.

Status

The current status of the LUN.

Thin

The disk format of the LUN, either thick or thin provisioned.

% Used/Free

A status bar that displays the used space on the LUN.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Dell EMC Unity Pools Report

Run the Dell EMC Unity Pools report to display the storage utilization of your Dell EMC Unity pools. View the total and free space available and the number of volumes and disks that make up your Dell EMC Unity pools.

BEFORE YOU BEGIN:

- Create and run a Dell EMC Unity Inventory job. You can select one or more Dell EMC Unity providers in a single job definition for cataloging. See [Create an Inventory Job Definition - Dell EMC Unity](#) on page 175.

Use the Dell EMC Unity Pools report to answer questions such as:

- What is the volume count of a specific storage pool?
- How many child pools are associated with a specific storage pool?

Parameters

Use the following parameters to customize your report:

- Storage Array
Multiple selections are supported.
- Only View Pools Exceeding Alert Threshold
Enable to view storage pools in which the usage exceeds the warning threshold.

Quick View

The Quick View section displays a pie chart of used and free space on your storage pools. Use the Dell EMC Unity Storage parameter to display storage pools on all storage systems or a specific storage system.

Note: The Quick View section is only modified through the Dell EMC Unity Host parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the Dell EMC Unity Pools report:

Dell EMC Unity Storage

The physical server where your files are stored.

Storage Pools

The number of storage pools on the Dell EMC Unity storage system.

Total Space

The total amount of storage that is assigned to the storage pool.

Used Space

The total storage allocated to volumes within the storage pool.

Available Space

The amount of storage that is assigned to the storage pool that is unused.

Detail View

The following fields and corresponding data display in the Detail View section of the DellEMC Unity Pools report:

Storage Pool

The name of the storage pool.

DellEMC Unity Storage

The physical server where your files are stored.

Total Space

The total amount of storage that is assigned to the storage pool.

Used Space

The total storage allocated to volumes within the storage pool.

Available Space

The amount of storage that is assigned to the storage pool that is unused.

Disks

The number of disks in the storage pool.

Status

The status of the disk with the highest priority status in the group.

Alert Threshold

A warning is generated when the assigned amount of space in the storage pool exceeds this level.

% Used/Free

A status bar that displays the used space on the storage pool.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

HPE Nimble Storage Volumes Report

Report the storage utilization of your HPE Nimble Storage volumes. Run the HPE Nimble Storage Volumes report to review the total capacity of your volumes, the total free space, and the percentage available to ascertain your HPE Nimble Storage volume storage utilization.

BEFORE YOU BEGIN:

- Create and run a HPE Nimble Storage Inventory job. You can select one or more HPE Nimble Storage providers in a single job definition for cataloging. See [Create an Inventory Job Definition - HPE Nimble Storage](#) on page 189.

Parameters

Use the following parameters to customize your report:

- Storage Array
Multiple selections are supported.

Quick View

The Quick View section displays the overall storage utilization of your HPE Nimble Storage volumes.

Summary View

The following fields and corresponding data display in the Summary View section of the HPE Nimble Storage Volumes report:

Storage Array

The physical storage system where your files are stored.

Total

The storage capacity that is available to the host.

Used

The used space on the storage array.

Empty Space

The available space on the storage array.

Snapshot Usage

The total capacity of the HPE Nimble Storage snapshots.

Volume Usage

The space allocated to HPE Nimble Storage volumes.

Shared Space

The total amount of shared space on the HPE Nimble Storage.

Total Reduction

The ratio of the total data reduced on the HPE Nimble Storage host. It includes data reduction, thin provisioning, zero detection, and unmap.

Data Reduction

The data reduction ratio of the HPE Nimble Storage. It includes deduplication, compression, and copy reduction.

% Used/Free

A status bar that displays the used space on the volume.

Detail View

The following fields and corresponding data display in the Detail View section of the IBM Spectrum Virtualize Storage Volumes report:

Volume

The name of the volume on the HPE Nimble Storage.

Storage Array

The physical storage system where your files are stored.

Provisioned

The total provisioned storage space on the HPE Nimble Storage volume.

Used

The used space on the volume.

Snapshot Usage

The total space on the HPE Nimble Storage occupied by snapshots.

Volume Usage

The total space on the HPE Nimble Storage used by volumes.

Total Reduction

The ratio of the total data reduced on the HPE Nimble Storage volume. It includes data reduction, thin provisioning, zero detection, and unmap.

Data Reduction

The data reduction ratio of the HPE Nimble Storage volume. It includes deduplication, compression, and copy reduction.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

IBM Spectrum Accelerate Pools Report

Run the IBM Spectrum Accelerate Pool report to review the storage utilization of your IBM Spectrum Accelerate Pools. View the total and free space available and the number of volumes and disks that make up your pools.

BEFORE YOU BEGIN:

- Create and run an IBM Spectrum Accelerate Inventory job. You can select one or more IBM providers in a single job definition for cataloging. See [Create an Inventory Job Definition - IBM Spectrum Accelerate](#) on page 177.

Parameters

Use the following parameters to customize your report:

- Storage Array
Multiple selections are supported.
- Only View Pools Exceeding Alert Threshold
Enable to view pools in which the usage exceeds the warning threshold.

Detail View

The following fields and corresponding data display in the Detail View section of the IBM Spectrum Virtualize Pools report:

Storage Pool

The name of the storage pool as known to IBM Spectrum Copy Data Management.

Storage Array

The name of the storage array as known to IBM Spectrum Copy Data Management.

Thin

The disk format of the storage system, either thick or thin provisioned.

Hard Usage

A status bar that displays the hard size, or physical capacity, of the storage pool.

Soft Usage

A status bar that displays the soft size, or maximum size seen by the hosts, of the storage pool.

Snapshot Usage

A status bar that displays the space used by snapshots on the volumes.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

IBM Spectrum Accelerate Volumes Report

Report the storage utilization of your IBM Spectrum Accelerate volumes. Run the IBM Spectrum Accelerate Volumes report to view the total space used by and free space available to your IBM Spectrum Accelerate volumes.

BEFORE YOU BEGIN:

- Create and run an IBM Spectrum Accelerate Inventory job. You can select one or more IBM providers in a single job definition for cataloging. See [Create an Inventory Job Definition - IBM Spectrum Accelerate](#) on page 177.

Parameters

Use the following parameters to customize your report:

- Storage Array
 - Multiple selections are supported.
- Detail View Filter

The default report parameters report on all storage arrays with more than 80% space used.

Quick View

The Quick View section displays the overall volume utilization of your IBM storage volumes.

Summary View

The following fields and corresponding data display in the Detail View section of the IBM Spectrum Accelerate Volumes report:

Storage Array

The name of the storage array as known to IBM Spectrum Copy Data Management.

of Pools

The number of storage pools on the storage array.

of Volumes

The number of volumes on the storage array.

Size

The storage capacity that is available to a host.

Used Capacity

The total used storage on the storage array.

Snapshots Used Capacity

The total used storage on the storage array dedicated to snapshots.

Detail View

The following fields and corresponding data display in the Detail View section of the IBM Spectrum Accelerate Storage Volumes report:

Volume

The name of the volume on the IBM storage system.

Storage Array

The physical storage system where your files are stored

Pool

The volume's associated storage pool.

Size

The volume storage capacity that is available to a host.

Used Capacity

The total used storage on the storage array.

Snapshots Used Capacity

The total used storage on the storage array dedicated to snapshots.

Consistency Group

The volume's associated consistency group.

Locked Status

The volume's lock status. If a volume is locked, no write commands are allowed.

% Used/Free

A status bar that displays the used space on the volume.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

IBM Spectrum Virtualize Consistency Groups Report

Run the IBM Spectrum Virtualize Consistency Group report to display information about your IBM Consistency Groups. View the associated storage providers, source and target volumes, and protection status of your IBM volumes through Consistency Groups.

BEFORE YOU BEGIN:

- Create and run an IBM Inventory job. You can select one or more IBM providers in a single job definition for cataloging. See [Create an Inventory Job Definition - IBM Spectrum Virtualize](#) on page 181 and [Create an Inventory Job Definition - IBM Spectrum Protect Snapshot](#) on page 179.

Use the IBM Consistency Group report to answer questions such as:

- What is the mapping name of a specific Consistency Group?
- What are the source and target volumes associated with a Consistency Group relationship?

Parameters

Use the following parameters to customize your report:

- Storage Array
Multiple selections are supported.
- Protection Type
Select FlashCopy or Global Mirror with Change Volumes protection type. Multiple selections are supported.

Consistency Groups - FlashCopy

The following fields and corresponding data display in the Consistency Groups - FlashCopy section of the IBM Spectrum Virtualize Consistency Groups report:

Consistency Group

The name of the Consistency Group along with its status (for example: copying or idle/copied) and FlashTime.

Storage Array

The name of the storage virtualizer as known to IBM Spectrum Copy Data Management where the consistency group resides.

Mapping Name

The mapping name IBM Spectrum Copy Data Management assigns to the Consistency Group, which consists of the Consistency Group name plus the source volume name.

Source Volume

The name of the source volume in the FlashCopy relationship.

Target Volume

The name of the target volume in the FlashCopy relationship.

Consistency Groups - Global Mirror with Change Volumes

The following fields and corresponding data display in the Consistency Groups - Global Mirror with Change Volumes section of the IBM Spectrum Virtualize Consistency Groups report:

Consistency Group

The name of the consistency group along with its status (for example: copying or idle/copied) and FlashTime, master and auxiliary cluster.

Storage Array

The name of the storage virtualizer as known to IBM Spectrum Copy Data Management where the consistency group resides.

Mapping Name

The mapping name IBM Spectrum Copy Data Management assigns to the Consistency Group.

Master Volume

The name of the master, or source, volume in the Global Mirror relationship.

Auxiliary Volume

The name of the auxiliary, or backup, volume in the Global Mirror relationship.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

IBM Spectrum Virtualize Pools Report

Run the IBM Spectrum Virtualize Pool report to display the storage utilization of your IBM storage pools. View the total and free space available and the number of volumes and disks that make up your IBM storage pools.

BEFORE YOU BEGIN:

- Create and run an IBM Spectrum Virtualize Inventory job. You can select one or more IBM providers in a single job definition for cataloging. See [Create an Inventory Job Definition - IBM Spectrum Virtualize](#) on page **181**.

Use the IBM Spectrum Virtualize Pools report to answer questions such as:

- What is the volume count of a specific storage pool?
- How many child pools are associated with a specific storage pool?

Parameters

Use the following parameters to customize your report:

- Storage Array
Multiple selections are supported.
- Detail View Filter
Available options include Warning Exceeded or All. Select Warning Exceeded to view storage pools in which the usage exceeds the warning threshold.

Quick View

The Quick View section displays a pie chart of used and free space on your storage pools. Use the IBM Host parameter to display storage pools on all storage systems or a specific storage system.

Note: The Quick View section is only modified through the IBM Host parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the IBM Spectrum Virtualize Pools report:

Storage Array

The name of the storage virtualizer as known to IBM Spectrum Copy Data Management.

Storage Pools

The number of storage pools on the IBM storage system.

Volume Count

The number of volumes that make up the storage pool.

Capacity

The total amount of MDisk storage that is assigned to the storage pool.

Allocated

The total storage allocated to volumes within the storage pool.

Virtual Capacity

The total virtual size of all the volume copies that are associated with the storage pool.

Child Pool Capacity

The capacity of the associated child pool, if available.

Child Pools

The number of associated child pools, if available.

Detail View

The following fields and corresponding data display in the Detail View section of the IBM Spectrum Virtualize Pools report:

Storage Pool

The name of the storage pool and associated child pools. Note that child pool capacities are not included in column totals.

Storage Array

The name of the storage virtualizer as known to IBM Spectrum Copy Data Management.

Capacity

The total amount of MDisk storage that is assigned to the storage pool.

Allocated

The total storage allocated to volumes within the storage pool.

Virtual Capacity

The total virtual size of all the volume copies that are associated with the storage pool.

External Virtual Capacity

The aggregate capacity of the managed and image mode MDisks from storage controllers virtualized using the "External Virtualization" feature of the chosen IBM storage systems.

Volumes

The number of volume copies that are in the storage pool.

Disk Count

The number of MDisks in the storage pool.

Status

The status of the MDisk with the highest priority status in the group, excluding image mode MDisks.

Warning

A warning is generated when the assigned amount of space in the storage pool exceeds this level.

%

The percentage of space used on the storage pool.

% Used/Free

A status bar that displays the used space on the storage pool.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

IBM Spectrum Virtualize Volumes Report

Report the storage utilization of your IBM storage volumes. Run the IBM Spectrum Virtualize Storage Volumes report to view the total capacity of your volumes, the total free space, and the available percentage.

BEFORE YOU BEGIN:

- Create and run an IBM Spectrum Virtualize Inventory job. You can select one or more IBM providers in a single job definition for cataloging. See [Create an Inventory Job Definition - IBM Spectrum Virtualize](#) on page **181**.

Use the IBM Spectrum Virtualize Storage Volumes report to answer questions such as:

- What is the total available storage space in the entire system?
- What is the amount of free and used space?
- How many volumes are available on a specific storage system?
- What is the size of the volume and the storage system that it resides on?

Parameters

Use the following parameters to customize your report:

- Storage Array
- Detail View Filter
- Show Only Thin Provisioned Volumes

Quick View

The Quick View section displays the overall volume utilization of your IBM storage volumes.

Summary View

The following fields and corresponding data display in the Summary View section of the IBM Spectrum Virtualize Storage Volumes report:

Storage Array

The name of the IBM storage system.

Volume Count

The number of volumes available on the IBM storage system.

Capacity

The volume storage capacity that is available to a host.

Real Capacity

The amount of physical storage that is allocated from a storage pool to volume copies.

Used Capacity

The portion of real capacity that is being used to store data.

Free Capacity

The difference between the real capacity and used capacity values for volume copies.

Detail View

The following fields and corresponding data display in the Detail View section of the IBM Spectrum Virtualize Storage Volumes report:

Volume

The name of the volume on the IBM storage system.

Storage Array

The physical storage system where your files are stored.

Storage Pool

The volume's associated storage pool.

Capacity

The difference between the real capacity and used capacity values for volume copies.

Real Capacity

The amount of physical storage that is allocated from a storage pool to volume copies.

Used Capacity

The portion of real capacity that is being used to store data.

Warning Threshold

For thin provisioned or compressed volume copies, a warning is generated at this percentage of the volume capacity.

Status

The status of the volume. A volume can be online, offline, or degraded.

Thin

Displays the thin provisioned status of the volume.

% Used/Free

A status bar that displays the used space on the volume.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Instant Disk Restore Volumes Report

Display a list of mapped volumes created through Instant Disk Restore jobs. View the name of the volume, location, and associated restore job.

BEFORE YOU BEGIN:

- Create and run an Instant Disk Restore job.

Parameters

Use the following parameters to customize your report:

- Storage Vendor

Set the storage vendor types to display in the report. Multiple selections are supported.

Detail View

The following fields and corresponding data display in the Detail View section of the Instant Disk Restore Volumes report:

Volume

The name of the mapped Instant Disk Restore volume on the node.

Location (Site)

The node and associated site where the volume resides.

Storage Vendor

The storage vendor of the node.

Job Name (Type)

The job and job type associated with the Instant Disk Restore.

Recovery Time

The time at which the restore job completed.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page 449
- [Report Overview](#) on page 360
- [Run a Report](#) on page 361
- [Create a Customized Report](#) on page 363

NetApp ONTAP Aggregates Report

Report the storage utilization and configuration of your aggregates before reaching your aggregate capacity. An aggregate is a collection of disks logically grouped together that provide storage to the volumes that they contain. Run the Aggregates report to view the total and free space available on your aggregate to help you determine if reallocation is necessary based on the size, the number of volumes and disks, and the percentage of space available on your aggregates.

BEFORE YOU BEGIN:

- Create and run a NetApp ONTAP Storage Inventory job. You can select one or more NetApp ONTAP providers in a single job definition for cataloging. See [Create an Inventory Job Definition - NetApp ONTAP Storage](#) on page 183.

Use the Aggregates report to answer questions such as:

- What is the total size of the aggregates?
- What is the amount of free and used space on the aggregate?
- How many disks make up the aggregate?

Parameters

Use the following parameters to customize your report:

- Storage Array
- Detail View Filter

The default report parameters report on all aggregates with more than 80% space used.

Quick View

The Quick View section displays a pie chart of used and free space on your aggregates. Use the NetApp ONTAP Storage parameter to display aggregates on all storage systems or a specific storage system.

Note: The Quick View section is only modified through the NetApp ONTAP Storage parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the Aggregates report:

Storage Array

The physical server where your files are stored.

Aggregate Count

The number of aggregates on the node.

Volume Count

The number of volumes that make up the aggregate.

Disk Count

The number of disks that make up the aggregate.

Total

The total size of the aggregate.

Available

The amount of free space available in the aggregate.

% Used

The percentage of used storage space on the aggregate.

Detail View

The following fields and corresponding data display in the Detail View section of the Aggregates report:

Aggregate

The name of the aggregate.

Storage Array

The physical server where your files are stored.

Location

The node where the volume resides.

Total

The total size of the aggregate.

Available

The amount of free space available in the aggregate.

Volume Count

The number of volumes that make up the aggregate.

Disk Count

The number of disks that make up the aggregate.

Status

The status of the aggregate. An aggregate can be online, offline for maintenance, or reserved for snapshot storage.

% Used/Free

A status bar that displays the used space on the aggregate.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

NetApp ONTAP LUNs Report

Review the storage utilization of your NetApp ONTAP LUNs. Run the NetApp ONTAP LUNs report to view your LUNs, total and available space, and online status.

BEFORE YOU BEGIN:

- Create and run a NetApp ONTAP Storage Inventory job. You can select one or more NetApp ONTAP providers in a single job definition for cataloging. See [Create an Inventory Job Definition - NetApp ONTAP Storage](#) on page 183.

Use the NetApp ONTAP LUNs report to answer questions such as:

- How many LUNs are associated with a storage system?
- What is the total and allocated capacity of a LUN?
- What LUNs are not being used, so I can reclaim this space?

Parameters

Use the following parameters to customize your report:

- Storage Array
 - Volume
- Multiple selections are supported.
- Detail View Filter

The default report parameters report on all LUNs with more than 80% space used.

Quick View

The Quick View section displays a pie chart of used and free space on your LUNs. Use the NetApp ONTAP Storage parameter to display LUNs on all storage systems or a specific storage system.

Note: The Quick View section is only modified through the NetApp ONTAP Storage parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the NetApp ONTAP LUNs report:

Storage Array

The physical server where your files are stored.

LUNs Count

The number of LUNs on the node.

Total Size

The total size of the aggregate.

Available Size

The amount of free space available in the aggregate.

% Used

The percentage of used storage space on the aggregate.

Detail View

The following fields and corresponding data display in the Detail View section of the NetApp ONTAP LUNs report:

LUN

The name of the LUN.

Storage Array

The physical server where your files are stored.

Volume

The volume associated with the LUN.

Location

The node where the volume resides.

Total Size

The total size of the LUN.

Available Size

The amount of free space available on the LUN.

Status

The status of the LUN. A LUN can be online or offline for maintenance.

Thin Provisioned

The disk format of the LUN, either thick or thin provisioned.

% Used/Free

A status bar that displays the used space on the LUN.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**

- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

NetApp ONTAP Orphaned LUNs Report

Review the NetApp ONTAP Orphaned LUNs Report to view LUNs that have no initiator group mapping or LUNs that belong to volumes that are offline.

BEFORE YOU BEGIN:

- Create and run a NetApp ONTAP Storage Inventory job. You can select one or more NetApp ONTAP providers in a single job definition for cataloging. See [Create an Inventory Job Definition - NetApp ONTAP Storage](#) on page 183.

Use the NetApp ONTAP Orphaned LUNs report to answer questions such as:

- How much space on an object is consumed by orphaned LUNs?
- Is thin provisioning enabled on a specific LUN?

Parameters

Use the following parameter to customize your report:

- Storage Array
- Volume

Multiple selections are supported.

The default report parameters report on all NetApp ONTAP storage volumes.

Quick View

The Quick View section displays a pie chart of space consumed by orphaned LUNs. Use the NetApp ONTAP Storage parameter to display orphaned LUNs on all storage systems or a specific storage system.

Note: The Quick View section is only modified through the NetApp ONTAP Storage parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the NetApp ONTAP Orphaned LUNs report:

Storage Array

The physical server where your files are stored.

Orphaned LUNs

The number of orphaned LUNs on the node.

Total Size (Volumes)

The total size of the volume on the node.

Total Size (Orphaned LUNs)

The total space on the volume occupied by orphaned LUNs.

% Used (Orphaned LUNs)

The percentage of used storage space on the node.

Detail View

The following fields and corresponding data display in the Detail View section of the NetApp ONTAP Orphaned LUNs report:

LUN

The name of the LUN.

Storage Array

The physical server where your files are stored.

Volume

The volume associated with the LUN.

Location

The node where the volume resides.

Thin Provisioned

The disk format of the LUN, either thick or thin provisioned.

Total Size

The total size of the LUN.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

NetApp ONTAP Quotas Report

Review the quotas of your NetApp ONTAP storage devices to determine which users or groups are approaching or exceeding their quota limits.

BEFORE YOU BEGIN:

- Create and run a NetApp ONTAP Storage Inventory job. You can select one or more NetApp ONTAP providers in a single job definition for cataloging. See [Create an Inventory Job Definition - NetApp ONTAP Storage](#) on page 183.

Parameters

Use the following parameters to customize your report:

- Storage Array
- Volume
 - Multiple selections are supported.
- Quota Criteria
- Top Quota Users

Detail View

The following fields and corresponding data display in the Detail View section of the NetApp ONTAP Quotas report:

Storage Array

The physical server where your files are stored.

Volume

The name of the volume on the node.

Location

The node where the volume resides.

Qtree

The name of the associated qtree.

Users

The users affected by the quota.

Quota Target

The name and location of the quota file on the volume.

Type

The type of entity to apply the quota against. For example, users, groups, or qtrees.

Space Usage

The space used on the volume.

Space Hard Limit

The hard disk space limit defined by the quota.

Space Soft Limit

The soft disk space limit defined by the quota.

% Used

The percentage of the quota used on the volume.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

NetApp ONTAP Snapshots Report

Review the storage utilization of your NetApp ONTAP Snapshots. Run the NetApp ONTAP Snapshots report to view the largest snapshots on your storage systems and the amount of space devoted to snapshot storage.

BEFORE YOU BEGIN:

- Create and run a NetApp ONTAP Storage Inventory job. You can select one or more NetApp ONTAP providers in a single job definition for cataloging. See [Create an Inventory Job Definition - NetApp ONTAP Storage](#) on page 183.

Use the NetApp ONTAP Snapshots report to answer questions such as:

- How many snapshots are on a storage system?
- What is the percentage of space on a volume that is used for snapshot storage?

Parameters

Use the following parameters to customize your report:

- Storage Array
- Volume
 - Multiple selections are supported.
- Number of Largest Snapshots to View

The default report parameters report on the hundred largest snapshots on all volumes.

Quick View

The Quick View section displays a pie chart of the size on your volumes consumed by snapshots. Use the NetApp ONTAP Storage parameter to display snapshots on all storage systems or a specific storage system.

Note: The Quick View section is only modified through the NetApp ONTAP Storage parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the NetApp ONTAP Snapshots report:

Storage Array

The physical server where your files are stored.

Snapshot Count

The number of snapshots on the node.

Total Volume Size

The total size of the volume on which the snapshots are stored.

Total Snapshot Size

The total combined size of all snapshots on the node.

% Used By Snapshot

The percentage of space on the volume used for snapshot storage.

Detail View

The following fields and corresponding data display in the Detail View section of the NetApp ONTAP Snapshots report:

Snapshot

The name of the snapshot.

Storage Array

The physical server where your files are stored.

Volume

The volume on which the snapshot is stored.

Location

The node where the volume resides.

Snapshot Creation Time

The snapshot creation date and time.

Volume Size

The total size of the volume on which the snapshot is stored.

Snapshot Size

The total size of the snapshot.

Total %

The percentage of space on the volume used by the snapshot.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

NetApp ONTAP Volumes Report

Manage your storage needs and review your volume storage capacity. Run the Volumes report to view the total capacity of your volumes, the total free space, and the available percentage.

BEFORE YOU BEGIN:

- Create and run a NetApp ONTAP Storage Inventory job. You can select one or more NetApp ONTAP providers in a single job definition for cataloging. See [Create an Inventory Job Definition - NetApp ONTAP Storage](#) on page 183.

Use the NetApp ONTAP Volumes report to answer questions such as:

- What is the total available storage space in the entire system?
- What is the amount of free and used space?
- How many volumes are available on a specific storage system?
- What is the size of the volume and the storage system that it resides on?

Parameters

Use the following parameters to customize your report:

- Storage Array
Multiple selections are supported.
- Detail View Filter

The default report parameters report on all NetApp ONTAP storage systems with more than 80% space used.

Quick View

The Quick View section displays a pie chart of used and free space on your volumes. Use the NetApp ONTAP Storage parameter to display volumes on all storage systems or a specific storage system.

Note: The Quick View section is only modified through the NetApp ONTAP Storage parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the NetApp ONTAP Volumes report:

Storage Array

The number of nodes included in the report, based on your parameters.

Volume Count

The number of cataloged volumes included in the report, based on your parameters.

Total

The total space of the volumes included in the report.

Available

The space available on the volumes included in the report.

Reserved

The space reserved for snapshot storage on the volume.

% Used

The percentage of used storage space on the volumes included in the report.

Detail View

The following fields and corresponding data display in the Detail View section of the NetApp ONTAP Volumes report:

Volume

The name of the volume on the node.

Storage Array

The physical server where your files are stored.

Aggregate

The name of the associated aggregate.

Location

The node where the volume resides.

Total

The total space on the volume.

Available

The space available on the volume.

Reserved

The space reserved for snapshot storage on the volume.

Status

The online status of the volume.

% Used / Free

The percentage of space used and a status bar that displays the used and free space on the volume. Note that space reserved for snapshot storage is included in this percentage.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Pure Storage FlashArray Volumes Report

Report the storage utilization of your Pure Storage FlashArray volumes. Run the Pure Storage FlashArray Volumes report to review the total capacity of your volumes, the total free space, and the percentage available to ascertain your Pure Storage FlashArray volume storage utilization.

BEFORE YOU BEGIN:

- Create and run a Pure Storage FlashArray Inventory job. You can select one or more Pure Storage providers in a single job definition for cataloging. See [Create an Inventory Job Definition - Pure Storage FlashArray](#) on page **192**.

Parameters

Use the following parameters to customize your report:

- Storage Array
Multiple selections are supported.

Quick View

The Quick View section displays the overall storage utilization of your Pure Storage FlashArray volumes.

Summary View

The following fields and corresponding data display in the Summary View section of the Pure Storage FlashArray Volumes report:

Storage Array

The physical storage system where your files are stored.

Total

The storage capacity that is available to the host.

Used

The used space on the storage array.

Empty Space

The available space on the storage array.

Snapshot Usage

The total capacity of the Pure Storage FlashArray snapshots.

Volume Usage

The space allocated to Pure Storage FlashArray volumes.

Shared Space

The total amount of shared space on the Pure Storage FlashArray.

Total Reduction

The ratio of the total data reduced on the Pure Storage FlashArray host. It includes data reduction, thin provisioning, zero detection, and unmap.

Data Reduction

The data reduction ratio of the Pure Storage FlashArray. It includes deduplication, compression, and copy reduction.

% Used/Free

A status bar that displays the used space on the volume.

Detail View

The following fields and corresponding data display in the Detail View section of the IBM Spectrum Virtualize Storage Volumes report:

Volume

The name of the volume on the Pure Storage FlashArray.

Storage Array

The physical storage system where your files are stored.

Provisioned

The total provisioned storage space on the Pure Storage FlashArray volume.

Used

The used space on the volume.

Snapshot Usage

The total space on the Pure Storage FlashArray occupied by snapshots.

Volume Usage

The total space on the Pure Storage FlashArray used by volumes.

Total Reduction

The ratio of the total data reduced on the Pure Storage FlashArray volume. It includes data reduction, thin provisioning, zero detection, and unmap.

Data Reduction

The data reduction ratio of the Pure Storage FlashArray volume. It includes deduplication, compression, and copy reduction.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

Storage Capacity Report

Report the storage capacity of your IBM Spectrum Virtualize/Accelerate pools, DellEMC Unity pools, Pure Storage FlashArray volumes, and NetApp ONTAP aggregates. Run the Storage Capacity report to view the total capacity of your volumes, the total free space, and the storage vendor.

BEFORE YOU BEGIN:

- Register storage providers, then create and run Inventory jobs. See [Register a Provider](#) on page 75 and [Jobs Overview](#) on page 161.
- Storage controllers will only be displayed if they have at least one resource successfully backed up to it.

Parameters

Use the following parameter to customize your report:

- Storage Vendor
- Show Managed Capacity Details

Enable to display a detailed view of the managed capacity of storage volumes within a storage array

Summary View

The following fields and corresponding data display in the Summary View section of the Storage Capacity report:

Storage Array

The name of the storage array.

Storage Vendor

The name of the storage vendor of the associated storage provider.

Usable Capacity

The total storage capacity that is available to a storage provider.

Managed Capacity

The space used by IBM Spectrum Copy Data Management backup jobs on the volumes of an array.

Detail View - Managed Capacity

The Detail View - Managed Capacity section displays if the Show Managed Capacity Details parameter is enabled. The Detail View displays the managed capacity of individual storage volumes within a storage array. The Replication field displays whether a replication relationship is associated with the volume. When calculating the Managed Capacity, the replication relationships are accounted for. If a replication relationship is available, the used size is doubled and displays in the Managed Capacity field.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

VM and Storage Mapping Report

Review the mappings of virtual machines down to the physical storage from which the datastore is created.

BEFORE YOU BEGIN:

- Create and run a VMware Inventory job. You can select one or more VMware providers in a single job definition for cataloging. See [Create an Inventory Job Definition - VMware](#) on page 194.

Parameters

Use the following parameters to customize your report:

- vCenter
- ESX Host
- NetApp ONTAP Storage
- Site

Detail View

The following fields and corresponding data display in the Detail View section of the VM Storage and Mapping report:

Datastore

The name of the virtual machine, host, and site and the name of the datastore that is used.

Disk

The disk on which the virtual machine is stored.

Path

The path to the virtual machine disk image file.

LUN

The corresponding logical unit number.

Source NetApp Node : Volume Used/Total

The corresponding NetApp ONTAP Node, volume, and the amount of used and total space.

Replication Destination

If a replication destination exists, it is listed here.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

VMware Datastores Report

Review the storage utilization of your VMware datastores, including the total free space, provisioned space, and capacities. Run the VMware Datastores report to view your datastores, the number of virtual machines on the datastores, and the percentage of space available.

BEFORE YOU BEGIN:

- Create and run a VMware Inventory job. You can select one or more VMware providers in a single job definition for cataloging. See [Create an Inventory Job Definition - VMware](#) on page 194.

Use the VMware Datastores report to answer questions such as:

- What is the file system type of a datastore?
- How many virtual machines are available on a datastore?
- What is the host node where the datastore resides?
- What is the free and provisioned space on a datastore?

Parameters

Use the following parameters to customize your report:

- vCenter
- ESX Host
 - Multiple selections are supported.
- Detail View Filter

The default report parameters report on all datastores with 80% space used, all vCenters, and all ESX Hosts.

Quick View

The Quick View section displays a pie chart of used and free space on your datastores. Use the ESX Host parameter to display datastores on all hosts or a specific host.

Note: The Quick View section is only modified through the ESX Host parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the VMware Datastores report:

Datastore Type

The file system types used by your datastores. For example, NFS or VMFS.

Datastore Count

The number of datastores associated with the datastore type.

Capacity

The total capacity of the datastore by file system type.

Provisioned Space

The amount of space on the datastore allocated for virtual disk files by file system type.

Free Space

The space available on the datastore by file system type.

Detail View

The following fields and corresponding data display in the Detail View section of the VMware Datastores report:

Datastore

The name of the datastore.

ESX Host (vCenter)

The host node where the datastore resides. More than one datastore can reside on an ESX host.

Type

The file system type of the datastore. For example, NFS or VMFS.

VM Count

The number of virtual machines on the datastore.

Capacity

The capacity of the datastore.

Provisioned Space

The amount of space on the datastore allocated for virtual disk files.

Free Space

The space available on the datastore.

% Used/Free

The percentage of space used on the datastore and a visual indicator of the amount of space used and available.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**

- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

VMware LUNs Report

Review the storage utilization of your VMware LUNs, including the amount of data transferred through the available transport types. Run the VMware LUNs report to view your LUNs, associated datastores, total and allocated capacities, and storage vendors.

BEFORE YOU BEGIN:

- Create and run a VMware Inventory job. You can select one or more VMware providers in a single job definition for cataloging. See [Create an Inventory Job Definition - VMware](#) on page 194.

Use the VMware LUNs report to answer questions such as:

- How many LUNs are associated with a specific storage vendor?
- What is the total and allocated capacity of a LUN?

Parameters

Use the following parameters to customize your report:

- vCenter
- ESX Host

Multiple selections are supported.

The default report parameters report on all vCenters and ESX hosts.

Quick View

The Quick View section displays a pie chart of LUN storage utilization by storage vendor. Use the ESX Host parameter to display LUNs on all hosts or a specific host. Any storage vendor that takes less than 5% of the total size of all datastores displays as Others.

Note: The Quick View section is only modified through the ESX Host parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the VMware LUNs report:

ESX Host (vCenter)

The host node where the LUN resides. More than one LUN can reside on an ESX host.

Fiber Channel

The capacity of the storage attached through fiber channel.

iSCSI

The capacity of the storage attached through iSCSI.

Block Adapter

The capacity of the storage attached through a block adapter.

Parallel SCSI

The capacity of the storage attached through parallel SCSI.

Detail View

The following fields and corresponding data display in the Detail View section of the VMware LUNs report:

LUN Name

The name of the LUN.

LUN ID

The unique identification number of the LUN.

Storage Vendor

The name of the storage vendor of the LUN.

ESX Host (vCenter)

The host node where the LUN resides.

Datastore(s)

The name of the associated datastore.

Capacity

The total capacity of the LUN.

Transport Type

The method through which data is transferred. For example, fiber channel or iSCSI.

RDM

The raw device mapping type. For example, physical or virtual.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

VMware Orphaned Datastores Report

Review the datastores that do not have any virtual machines assigned to them, or if virtual machines are assigned to the datastores, view the virtual machines that are in an inaccessible state.

BEFORE YOU BEGIN:

- Create and run a VMware Inventory job. You can select one or more VMware providers in a single job definition for cataloging. See [Create an Inventory Job Definition - VMware](#) on page 194.

Parameters

Use the following parameter to customize your report:

- vCenter
- ESX Host

Multiple selections are supported.

The default report parameters report on all ESX hosts and vCenters.

Orphaned Datastores

The following fields and corresponding data display in the Orphaned Datastores section of the VMware Orphaned Datastores report.

Datastore

The name of the datastore.

ESX Host (vCenter)

The host node where the datastore resides. More than one datastore can reside on an ESX host.

Type

The file system type of the datastore. For example, NFS or VMFS.

Capacity

The capacity of the datastore.

Provisioned Space

The amount of space on the datastore allocated for virtual disk files.

Free Space

The amount of free space on the datastore.

% Used/Free

The percentage of space used on the datastore and a visual indicator of the amount of space used and available.

Reason

The reason the datastore was returned by the report. For example, if no virtual machines are registered on the datastore.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

VMware Orphaned LUNs Report

Review VMware orphaned LUNs. These are the LUNs that are neither used as datastores nor RDMs.

BEFORE YOU BEGIN:

- Create and run a VMware Inventory job. You can select one or more VMware providers in a single job definition for cataloging. See [Create an Inventory Job Definition - VMware](#) on page 194.

Use the VMware Orphaned LUNs report to answer questions such as:

- What is the transport type of an orphaned LUN?
- What is the storage vendor of an orphaned LUN?

Parameters

Use the following parameter to customize your report:

- vCenter
- ESX Host

Multiple selections are supported.

The default report parameters report on all ESX hosts and vCenters.

Summary View

The following fields and corresponding data display in the Detail View section of the Orphaned LUNs report.

vCenter

The name of the vCenter node.

ESX Host

The host node where LUNs reside. More than one LUN can reside on an ESX host.

Total LUNs

The total number of LUNs on the host.

Capacity

The capacity overall LUN storage utilization.

Detail View

The following fields and corresponding data display in the Detail View section of the Orphaned LUNs report.

LUN Name

The name of the LUN.

LUN ID

The unique identification number of the LUN.

Storage Vendor

The name of the storage vendor of the LUN.

ESX Host (vCenter)

The host node where the LUN resides. More than one LUN can reside on an ESX host.

Transport Type

The method through which data is transferred. For example, fibre channel or iSCSI.

Capacity

The capacity of the LUN.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

VMware VM Snapshot Sprawl Report

The VMware VM Snapshot Sprawl report displays the age and number of snapshots used to protect your virtual machines through NetApp ONTAP Snapshot software.

BEFORE YOU BEGIN:

- Create and run a VMware Inventory job. You can select one or more VMware providers in a single job definition for cataloging. See [Create an Inventory Job Definition - VMware](#) on page 194.

Use the VMware VM Snapshot Sprawl report to answer questions such as:

- What is the age of the oldest snapshot on a virtual machine?
- Which virtual machines have a large number of snapshots?

Parameters

Use the following parameters to customize your report:

- vCenter
- ESX Host
 - Multiple selections are supported.
- Snapshot Sprawl Criteria
- Snapshot Creation Time

The default report parameters report on all the criteria with a Snapshot Creation time of more than a year.

Detail View - VMs with Aged/Memory Snapshots

The following fields and corresponding data display in the Detail View - VMs with Aged Snapshots and VMs with Memory Snapshots:

VM Name

The name of the virtual machine along with the location of the host node where the virtual machine resides.

Snapshot Name

The name of the oldest snapshot on the virtual machine.

Snapshot Creation Time

The creation date and time of the oldest snapshot on the virtual machine.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

VMware VM Sprawl Report

Review the status of your virtual machines, including virtual machines that are powered off, powered on, or suspended. Run the VMware VM Sprawl report to view unused virtual machines, the date and time they were powered off, and virtual machine templates.

BEFORE YOU BEGIN:

- Create and run a VMware Inventory job. You can select one or more VMware providers in a single job definition for cataloging. See [Create an Inventory Job Definition - VMware](#) on page 194.

Use the VMware VM Sprawl report to answer questions such as:

- What are the names of my virtual machines and associated datastores?
- How many virtual machines have been powered off for more than 30 days, 90 days, 180 days, or one year?
Can I reclaim this space?

Parameters

Use the following parameters to customize your report:

- vCenter
- ESX Host
Multiple selections are supported.
- Days Since Last Power Off
- Days Since Last Suspended
- Days Since Last Power On

The default report parameters report on all criteria that were last powered on over 180 days ago.

Quick View

The Quick View section displays a pie chart of used and free space on your virtual machines. Use the ESX Host parameter to display virtual machines on all hosts or a specific host.

Note: The Quick View section is only modified through the ESX Host parameter.

Detail View - Powered Off VMs

The following fields and corresponding data display in the Detail View - Powered Off VMs section of the VMware VM Sprawl report:

VM Name

The name of the virtual machine.

Powered Off Since

The date and time the virtual machine was last powered off.

ESX Host (vCenter)

The host node where the virtual machine resides.

Resource Pool

The name of the associated resource pool.

Provisioned Space

The amount of space on the datastore allocated for virtual disk files.

Datastore(s)

The name of the associated datastores.

Detail View - Suspended VMs

The following fields and corresponding data display in the Detail View - Suspended VMs section of the VMware VM Sprawl report:

VM Name

The name of the virtual machine.

Suspended Since

The date and time elapsed since the virtual machine was suspended.

ESX Host (vCenter)

The host node where the virtual machine resides. More than one datastore can reside on an ESX host.

Resource Pool

The name of the associated resource pool.

Provisioned Space

The amount of space on the datastore allocated for virtual disk files.

Datastore(s)

The name of the associated datastores.

Detail View - Templates

The following fields and corresponding data display in the Detail View - Templates section of the VMware VM Sprawl report:

Template Name

The name of the virtual machine template.

ESX Host (vCenter)

The host node where the virtual machine template resides.

Provisioned Space

The amount of space on the datastore allocated for virtual disk files.

Datastore(s)

The name of the associated datastores.

Detail View - Powered On VMs

The following fields and corresponding data display in the Detail View - Powered On VMs section of the VMware VM Sprawl report:

VM Name

The name of the virtual machine.

Powered On Since

The date and time the virtual machine was last powered on.

ESX Host (vCenter)

The host node where the virtual machine resides.

Resource Pool

The name of the associated resource pool.

Provisioned Space

The amount of space on the datastore allocated for virtual disk files.

Datastore(s)

The name of the associated datastores.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page **449**
- [Report Overview](#) on page **360**
- [Run a Report](#) on page **361**
- [Create a Customized Report](#) on page **363**

VMware VM Storage Report

Review your virtual machines and associated datastores through the VMware VM Storage report. View datastores, resource pools, and the provisioned space of the datastore.

BEFORE YOU BEGIN:

- Create and run a VMware Inventory job. You can select one or more VMware providers in a single job definition for cataloging. See [Create an Inventory Job Definition - VMware](#) on page 194.

Parameters

Use the following parameters to customize your report:

- vCenter
- ESX Host

Detail View

The following fields and corresponding data display in the Detail View section of the VMware VM Storage report:

VM

The name of the virtual machine, along with the associated datastore.

ESX Host (vCenter)

The host node where the virtual machine resides.

Resource Pool

The name of the associated resource pool.

Provisioned Space

The amount of space on the datastore allocated for virtual disk files.

RELATED TOPICS:

- [Storage Utilization Reports](#) on page 449
- [Report Overview](#) on page 360
- [Run a Report](#) on page 361
- [Create a Customized Report](#) on page 363

Maintenance

The topics in the following section cover maintenance information including updating the IBM Spectrum Copy Data Management appliance, logging on to the virtual appliance, and collecting logs for troubleshooting.

Maintenance Overview

In most cases, IBM Spectrum Copy Data Management is installed on a virtual appliance. The virtual appliance contains the application and the Inventory.

System Administrators can perform maintenance tasks on the IBM Spectrum Copy Data Management application. Note that a System Administrator is usually a senior-level user who designed or implemented the vSphere and ESX infrastructure, or a user with an understanding of IBM Spectrum Copy Data Management, VMware, and Linux command-line usage. Maintenance tasks are performed in vSphere Client, through the IBM Spectrum Copy Data Management command-line, or through a web-based management console.

Maintenance tasks include collecting logs, updating the application, and reviewing the configuration of the virtual appliance.

RELATED TOPICS:

- [Log On to the Virtual Appliance](#) on page **515**
- [Set Time Zone](#) on page **516**
- [Collect Logs For Troubleshooting](#) on page **518**
- [Modifying Job Log Options](#) on page **521**
- [Manage the Administrative Console](#) on page **523**
- [Update IBM Spectrum Copy Data Management](#) on page **524**
- [Install the Marketplace RPM](#) on page **527**
- [Modifying Network Settings](#) on page **530**
- [Upload an SSL Certificate](#) on page **532**
- [Horizontal Scale Out](#) on page **534**
- [Restoring A Snapshot From A FlexGroup Volume To Another FlexGroup Volume](#) on page **539**
- [LDAP User Name Syntax](#) on page **552**

Log On to the Virtual Appliance

Log on to the IBM Spectrum Copy Data Management virtual appliance through vSphere Client to access the command prompt. Available options include collecting logs for troubleshooting.

BEFORE YOU BEGIN:

- Install IBM Spectrum Copy Data Management by deploying an OVF template. This creates a virtual machine containing the application on a VMware host such as an ESX or ESXi Server. See [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page 63.

To access the virtual appliance command prompt:

1. In vSphere Client, select the virtual machine where IBM Spectrum Copy Data Management is deployed.
2. In the **Summary** tab, select **Open Console** and click in the console.
3. Select **Login**, and enter your user name and password. The default user name is **administrator** and the default password is **ecxadLG235**.

To log off, enter **exit**.

RELATED TOPICS:

- [Collect Logs For Troubleshooting](#) on page 518
- [Manage the Administrative Console](#) on page 523

Set Time Zone

Access the IBM Spectrum Copy Data Management Administrative Console to set the time zone of the IBM Spectrum Copy Data Management appliance.

BEFORE YOU BEGIN:

- Install IBM Spectrum Copy Data Management by deploying an OVF template. This creates a virtual machine containing the application on a VMware host such as an ESX or ESXi Server. See [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page 63.

To set a time zone:

1. From a supported browser, enter the following URL:
https://<HOSTNAME>:8090/
where <HOSTNAME> is the IP address of the virtual machine where the application is deployed.
2. In the login window, select **System** from the **Authentication Type** drop-down menu. Enter your password to access the Administrative Console. The default password is **ecxadLG235**.
3. Click **Perform System Actions**.
4. In the Change Time zone section, select your time zone. A message stating the operation was successful displays. All IBM Spectrum Copy Data Management logs and schedules will reflect the selected time zone. The selected time zone will also display on the IBM Spectrum Copy Data Management appliance when logged in as a root user.
5. To view the current time zone, select **Product Information** from the main Administrative Console page. The displayed time updates every five seconds.

If instances of IBM Spectrum Copy Data Management reside in time zones that differ from the IBM Spectrum Copy Data Management appliance, the instances can be synced to the time zone of the IBM Spectrum Copy Data Management appliance instead of the time zone of the instance. In some cases, a user may wish to schedule IBM Spectrum Copy Data Management jobs based on the IBM Spectrum Copy Data Management appliance's time zone over their local time zone. Perform the following procedure to sync the time zones of the IBM Spectrum Copy Data Management instances with the IBM Spectrum Copy Data Management appliance:

To sync time zones of IBM Spectrum Copy Data Management instances with the IBM Spectrum Copy Data Management appliance:

1. Follow the procedure above to set the time zone of the IBM Spectrum Copy Data Management appliance.
2. Log in to the IBM Spectrum Copy Data Management appliance as a root user, and edit the following file: `/opt/virgo/repository/ecx-usr/com.syncsort.dp.xsb.api.endeavour.session.properties`.

3. In the property file, set `useServerTime` to `true`, then save the file.
4. Restart the IBM Spectrum Copy Data Management appliance through the following commands:

```
service virgo stop  
service virgo start
```

RELATED TOPICS:

- [Log On to the Virtual Appliance](#) on page **515**
- [Manage the Administrative Console](#) on page **523**

Collect Logs For Troubleshooting

For troubleshooting the IBM Spectrum Copy Data Management application, IBM Spectrum Copy Data Management can generate an archive of logs containing various files.

There are two approaches for downloading logs. Download logs from the Support menu or access the IBM Spectrum Copy Data Management virtual appliance through vSphere Client to download logs using a command prompt. The first approach is simpler and generally sufficient. The second approach produces a more comprehensive set of logs.

Audit logs can also be viewed and downloaded through the **Support**  menu. The Audit Log window displays a log of actions performed in IBM Spectrum Copy Data Management. Information included in the logs include the access time in EPOCH format, IBM Spectrum Copy Data Management appliance time, IP address of the requester, the username of the requester, the LDAP group to which the user belongs, the operation being performed, and readable text description of the action being performed.

Operations may be one of the following: create, read, update, and delete (CRUD).

Audit Logs can be searched and filtered based on the description, operation, time, or the username of the requester. The filtered log results can be downloaded for future archiving or an external audit.

BEFORE YOU BEGIN:

- Contact Technical Support to determine if they need a log collection file for troubleshooting.
- If you are collecting logs from the virtual appliance, ensure you have administrator access to the virtual appliance where IBM Spectrum Copy Data Management is deployed. See [User Administration and Security Management](#) on page 14.
- If you are collecting logs from the virtual appliance, download an SCP client to save the logs to your local machine. If your local computer is Windows based, you can use WinSCP. See <http://winscp.net/eng/index.php>. If your local computer is Unix based, you can use scp. See http://www.hypexr.org/linux_scp_help.php.

To collect audit logs from the Support menu:

1. Click the arrow next to the **Support**  icon, then click **View Audit Logs**.
2. The Audit Log window displays a log of actions performed in IBM Spectrum Copy Data Management, along with the user performing the action and a description of the action.
3. To search for the actions of a specific IBM Spectrum Copy Data Management user, search for the user name in the **Search for users** field.

4. To download the current view of the Audit Log as a .csv file, click **Download**, then select a location to save the file.

To collect logs from the Support menu:

1. Click the arrow next to the **Support**  icon, then click **Download Log Files** .
2. Select a location to save the zip file.

Note: The following logs are added to the zip file and saved to your local machine: mongo, rabbitmq, and virgo.

To collect IBM Spectrum Copy Data Management logs from the virtual appliance:

Note: This procedure assumes IBM Spectrum Copy Data Management deployment was to a VMware appliance host.

1. Log on to the virtual appliance console as **administrator**.

To log on to the virtual appliance:

1. In vSphere Client, select the virtual machine where IBM Spectrum Copy Data Management is deployed.
 2. In the **Summary** tab, select **Open Console** and click in the console.
 3. Select **Login**, and enter your user name and password. The default user name is **administrator** and the default password is **ecxadLG235**.
2. Navigate to **/opt/CDM/tools/scripts**.

Run **logcollect** using the **sudo** command. For example, at the command prompt enter:

```
$ sudo ./logcollect
```

or

```
$ sudo /opt/CDM/tools/scripts/logcollect
```

The logcollect script might take a few minutes to run depending on application usage. .

3. Optionally, add a specific job log to the archive using the **-job** command and the job ID, which can be obtained through the job's instance on the **Jobs**  tab. For example, at the command prompt enter:

```
$ sudo ./logcollect -job <job ID>
```

or

```
$ sudo /opt/CDM/tools/scripts/logcollect -job <job ID>
```

4. Find the collected logs archived as a zip file in **/opt/CDM/tools/logcollect/archives**. The zip file name is unique, created using the host or DNS name and timestamp. It is of the form **<host>_logs_<date>_<time>.zip**.

Note: The following logs are added to the zip file: logcollect, mongo, postgres, rabbitmq, system, virgo, and job if the **-job** command was run.

NEXT STEPS:

- Contact Technical Support to inform them that you have created a log collection file for troubleshooting.
- If you collected logs from the virtual appliance, copy the zip file to your local computer. If your local computer is Windows based, you can use WinSCP. See <http://winscp.net/eng/index.php>. If your local computer is Unix based, you can use scp. See http://www.hypexr.org/linux_scp_help.php.
- Send the zipped log collection file to Technical Support.
- Manually clean up the archive directory.

RELATED TOPICS:

- [Log On to the Virtual Appliance](#) on page **515**
- [Monitor a Job Session](#) on page **164**

Modifying Job Log Options

The settings for job logs can be edited for the IBM Spectrum Copy Data Management appliance.

BEFORE YOU BEGIN:

The settings used to control the job logs are listed below with their default values.

The `maintenance.joblog.enable` setting is a Boolean value that enables or disables the logging of jobs. The default setting for this is **false**.

The `maintenance.joblog.offload.enable` setting is a Boolean value that enables or disables the offload option for job logs. The default setting for this is **false**.

The `maintenance.joblog.maxrunningminutes` defines the maximum running time, in minutes, that a job log purge operation will run during a maintenance job. The default setting for this is **60** minutes.

`Maintenance.joblog.retentiondays` is the number of days that job logs are retained. The default setting for this is **30** days.

The `maintenance.joblog.export.target` controls the path to which job logs are exported and will only be used if the `maintenance.joblog.offload.enable` option is **true**. The default path for this is **/data2/joblogs**.

To modify job log options:

1. Log in to the IBM Spectrum Copy Data Management appliance as the root user.
2. Stop the virgo service on the IBM Spectrum Copy Data Management appliance.

```
$ service virgo stop
```
3. Using a text editor (vi), open the `com.syncsort.dp.xsb.serviceprovider.properties` file which is located at `/opt/virgo/repository/ecx-usr/`.
4. Modifying the file as necessary. The settings are listed with their defaults.

```
maintenance.joblog.enable=true  
maintenance.joblog.offload.enable=false  
maintenance.joblog.maxrunningminutes=60  
maintenance.joblog.retentiondays=30  
maintenance.joblog.export.target=/data2/joblogs
```
5. Save the properties file and exit the text editor.

6. Start the virgo service on the IBM Spectrum Copy Data Management appliance.

```
$ service virgo start
```

Manage the Administrative Console

Log on to the Administrative Console to review the configuration of the IBM Spectrum Copy Data Management virtual appliance. Available information includes general system settings, network, proxy settings, and available updates.

BEFORE YOU BEGIN:

- Install IBM Spectrum Copy Data Management by deploying an OVF template. This creates a virtual machine containing the application on a VMware host such as an ESX or ESXi Server. See [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page **63**.

CONSIDERATIONS:

-

To manage the Administrative Console:

1. From a supported browser, enter the following URL:
https://<HOSTNAME>:8090/
where <HOSTNAME> is the IP address of the virtual machine where the application is deployed.
2. In the login window, select **System** from the **Authentication Type** drop-down menu. Enter your password to access the Administrative Console. The default password is **ecxadLG235**.
3. Review the available options for the virtual appliance.

RELATED TOPICS:

- [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page **63**
- [Update IBM Spectrum Copy Data Management](#) on page **524**
- [Install the Marketplace RPM](#) on page **527**

Update IBM Spectrum Copy Data Management

BEFORE YOU BEGIN:

- Ensure IBM Spectrum Copy Data Management is installed and running. See [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page 63.
- Note that after IBM Spectrum Copy Data Management updates, it cannot rollback to a previous version without a virtual machine snapshot. Create a virtual machine snapshot of your environment before updating, then, if necessary, perform a virtual machine snapshot rollback to return to a previous version of IBM Spectrum Copy Data Management.

Note: If updating from IBM Spectrum Copy Data Management 2.2.5, or updating to a newer version of CDM that was previously updated from CDM 2.2.5, additional steps are required. Updates in this scenario will fail with the following error: `<open file '<fdopen>', mode 'rb' at (exception)>`. To resolve this issue, log into the CDM appliance with root access, then navigate to `/etc/yum.repos.d/`. If the `erlang_solutions.repo` file is in this location, move the file to `/tmp` on the appliance before updating. Once the file has been moved, run the update procedure again.

When updating IBM Spectrum Copy Data Management through the Administrative Console, a Clean Process Failure message displays during the update process. The message indicates that the update has been applied successfully, but the update cleanup process failed. This message appears consistently when performing an update to IBM Spectrum Copy Data Management, regardless of whether the Clean Process failed or was successful. This issue is resolved in the latest update.

To update your IBM Spectrum Copy Data Management appliance:

1. From a machine with internet access, download the update file from IBM Fix Central.
2. From a supported web browser, access the Administrative Console at the following address:
`https://<HOSTNAME>:8090/`
where `<HOSTNAME>` is the IP address of the virtual machine where the application is deployed.
3. In the login window, select **System** from the **Authentication Type** drop-down menu. Enter your password to access the Administrative Console. The default password is **ecxadLG235**.
4. Click **Manage updates**. In the **Don't want to update IBM Spectrum Copy Data Management to the latest available release?** section, select **Click here**.
5. Click **Choose File**, browse for the update file to upload to the appliance, then click **Upload Update Image**. The update process begins once the update image has been uploaded to the appliance.
6. After the update completes, navigate to the **Perform System Actions** page on the Administrative Console to restart the appliance.

HTML content from previous versions of IBM Spectrum Copy Data Management may be stored in your browser's cache. Clear your browser's cache before logging in to an updated version of IBM Spectrum Copy Data Management to ensure you are viewing the latest content changes.

Manually Update Administrative Console RPMs When Upgrading from IBM Spectrum Copy Data Management 2.2.6 to 2.2.7

When upgrading IBM Spectrum Copy Data Management from version 2.2.6 to 2.2.7 in an environment without internet access, additional steps must be taken to upgrade the Administrative Console. In summary, first the IBM Spectrum Copy Data Management appliance is updated to version 2.2.7 through an ISO, then three RPM packages are manually installed on the appliance.

1. From a machine with internet access, download the IBM Spectrum Copy Data Management 2.2.7 ISO from IBM Fix Central.
2. Contact Technical Support for information about accessing the Administrative Console RPM packages: `scdm-adminconsole-2.0.*`, `scdm-emi-2.0.*`, and `scdm-catalogmanager-1.0.*`.
3. Log in the 2.2.6 Administrative Console. Click **Manage updates**. In the Don't want to update to the latest available release? section, select **Click here**.
4. Click **Choose File**, browse for the 2.2.7 ISO to upload to the appliance, then click **Upload Update Image**. The update process begins once the update image has been uploaded to the appliance.
Once complete, the IBM Spectrum Copy Data Management appliance is updated to version 2.2.7.
5. Log out of the Administrative Console.
6. Log in to the IBM Spectrum Copy Data Management command line interface via vSphere or ssh utility as a root user.
7. Copy the Administrative Console RPM packages (`scdm-adminconsole-2.0.*`, `scdm-emi-2.0.*`, and `scdm-catalogmanager-1.0.*`) to the appliance's "/" directory.
8. Run the following commands as root:

```
rpm -Uvh scdm-adminconsole-2.0.0-*.rpm  
rpm -Uvh scdm-emi-2.0.0-*.rpm  
rpm -Uvh scdm-catalogmanager-1.0.0*.rpm
```
9. Once complete, clear your browser cache, then log into the updated IBM Spectrum Copy Data Management 2.2.7 Administrative Console.
10. Click **Product Information**, then verify that the RPM packages have been updated.

To update your IBM Spectrum Copy Data Management appliance through RPM update files:

The following procedure requires root access to the IBM Spectrum Copy Data Management virtual appliance.

1. From a machine with Internet Access, download the necessary RPM update files from IBM Fix Central.
2. Log into the IBM Spectrum Copy Data Management appliance as a root user.

3. Copy the RPM update files to the appliance's "/" directory.
4. Navigate to "/" through the cd command, then execute the following command for each RPM:

```
rpm -Uvh <rpm>
```
5. Once complete, restart the appliance.

NEXT STEPS:

- If the update process fails, encounters an error, or there is any interruption in the update process prior to it completing, review the update logs at the following location on the virtual machine: */opt/vmware/var/log/vami/updatecli.log*.
- To reapply the update, revert the virtual machine snapshot, which was created before the update procedure, and then attempt another update.

RELATED TOPICS:

- [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page **63**
- [Manage the Administrative Console](#) on page **523**

Install the Marketplace RPM

The Marketplace provides new IBM Spectrum Copy Data Management features and assets to add to your environment as they become available. Through Marketplace you can find and upload tools, utilities, reports, and updates.

BEFORE YOU BEGIN:

- Ensure IBM Spectrum Copy Data Management is installed and running. See [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page 63.

To install or update the Marketplace for IBM Spectrum Copy Data Management:

1. From a machine with Internet Access, download the **marketplaceadmin-<version>.noarch.rpm** file from IBM Fix Central at <https://www-945.ibm.com/support/fixcentral>
2. Copy **marketplaceadmin-<version>.noarch.rpm** to the IBM Spectrum Copy Data Management appliance "/" directory via SCP.
3. Log in as "root" user, and execute "cd /" to move to the "/" directory.
4. Execute one of the following commands:
 - To install the RPM package if it is not currently installed, execute "rpm -ivh marketplaceadmin-<version>.noarch.rpm"
 - To update an installed RPM package, execute "rpm -Uvh marketplaceadmin-<version>.noarch.rpm"

To upload a Marketplace asset:

1. Once RPM installation is done, access the Marketplace User Interface using your browser at <https://<ipaddress>:8443/mps/marketplace>. Log in using your IBM Spectrum Copy Data Management portal password.
2. Select desired asset to upload.

RELATED TOPICS:

- [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page 63
- [Manage the Administrative Console](#) on page 523

Backup and Restore the Catalog

The Catalog Manager provides IBM Spectrum Copy Data Management administrators with the ability to backup and restore the IBM Spectrum Copy Data Management Catalog. The Catalog Manager is available through the Administrative Console, which is available at <https://<HOSTNAME>:8090/>, where <HOSTNAME> is the IP address of the virtual machine where IBM Spectrum Copy Data Management is deployed.

BEFORE YOU BEGIN:

- Install IBM Spectrum Copy Data Management by deploying an OVF template. This creates a virtual machine containing the application on a VMware host such as an ESX or ESXi Server. See [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page 63.

To manage your IBM Spectrum Copy Data Management catalog through the Catalog Manager:

1. From a supported browser, enter the following URL:

https://<HOSTNAME>:8090/

where <HOSTNAME> is the IP address of the virtual machine where the application is deployed.

2. In the login window, select **System** from the **Authentication Type** drop-down menu. Enter your password to access the Administrative Console. The default password is **ecxadLG235**.
3. Click **Menu**, then select **Catalog Manager**.
4. Select Backup Catalog or Restore Catalog.

Backup Catalog: In the **Directory** field, enter the backup destination on the IBM Spectrum Copy Data Management host. Ensure the destination volume exists and that there is enough room on the destination volume for the Catalog backup. Click **Backup** to begin the Catalog backup.

Backup Catalog considerations: IBM Spectrum Copy Data Management will be stopped while the Catalog is being backed up. The IBM Spectrum Copy Data Management user interface will not be accessible, and all running jobs will be aborted.

Restore Catalog: In the **Directory** field, enter the source of the Catalog restore on the IBM Spectrum Copy Data Management host. Ensure the source volume exists. Click **Restore** to begin the Catalog restore.

Restore Catalog considerations: IBM Spectrum Copy Data Management will be stopped while the Catalog is being restore. The IBM Spectrum Copy Data Management user interface will not be accessible, and all running jobs will be aborted. All IBM Spectrum Copy Data Management snapshots created after the Catalog backup was run will be lost.

RELATED TOPICS:

- [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page **63**
- [Manage the Administrative Console](#) on page **523**

Modifying Network Settings

Access the IBM Spectrum Copy Data Management appliance using SSH to make adjustments to the network.

BEFORE YOU BEGIN:

- Install IBM Spectrum Copy Data Management by deploying an OVF template. This creates a virtual machine containing the application on a VMware host such as an ESX or ESXi Server. See [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page 63.

To modify the network settings:

Users can configure the network post deployment should it be required. Follow the process below to utilize the `system-config-network` tool. This tool is analogous to the NetworkManager Text User Interface (`nmtui`) for RHEL 7/CentOS 7. When connecting to the IBM Spectrum Copy Data Management appliance to use this tool, an account with root privileges will be required.

1. SSH from a terminal to the IBM Spectrum Copy Data Management appliance:

```
ssh root@appliancehostname
```

2. Enter password for account associated with the username.
3. Once successfully connected, enter the following to launch the tool:

```
system-config-network
```

Once changes have been made to the IP or hostname using the `system-config-network` tool, a reboot of the appliance may be required. Updates to any associated VADP proxy servers may be required.

To update VADP proxy server settings after an IBM Spectrum Copy Data Management appliance IP or hostname change:

Additionally, changes made to the IP or hostname of the IBM Spectrum Copy Data Management appliance will result in a loss of communication with associated VADP proxy servers. Please follow the steps below on each associated VADP proxy server to re-establish communication with the IBM Spectrum Copy Data Management appliance.

SSH to each VADP proxy and enter the following commands:

1. Enter the new IBM Spectrum Copy Data Management appliance IP or hostname to the `ECX_HOST` variable in `/opt/ECX/bin/escvadp` and save:

```
vi /opt/ECX/bin/ecxvadp  
ECX_HOST=<new ECX IP or Hostname>
```

2. Enter the new IBM Spectrum Copy Data Management appliance IP or hostname to the ECX_HOST variable in /etc/rd.d/init.d/escvadp and save:

```
vi /etc/rc.d/init.d/ecxvadp  
ECX_HOST=<new ECX IP or Hostname>
```

3. Restart the `ecxvadp` service:

```
service ecxvadp restart
```

Any associated VADP proxy servers should be now able to communicate with the IBM Spectrum Copy Data Management appliance using the updated IP address or hostname.

RELATED TOPICS:

- [Log On to the Virtual Appliance](#) on page **515**
- [Manage the Administrative Console](#) on page **523**

Upload an SSL Certificate

To establish secure connections in IBM Spectrum Copy Data Management, you must upload an SSL certificate through the web-based management console of the virtual machine where IBM Spectrum Copy Data Management is deployed.

BEFORE YOU BEGIN:

- Ensure IBM Spectrum Copy Data Management is installed and running. See [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page **63**.
- If you are uploading an LDAP SSL certificate, ensure an LDAP server is running and reachable by IBM Spectrum Copy Data Management.
- If you are uploading an LDAP SSL certificate, review LDAP syntax. See [LDAP User Name Syntax](#) on page **552**.

To upload a certificate:

1. Contact your network administrator for the name of the certificate to export.
2. From a supported browser, export the certificate to your computer. Make note of the location of the certificate on your computer. The process of exporting certificates varies based on your browser. See [Related Topics](#).
3. From a supported browser, enter the following URL:
https://<HOSTNAME>:8090/
where *<HOSTNAME>* is the IP address of the virtual machine where the application is deployed.
4. In the login window, select **System** from the **Authentication Type** drop-down menu. Enter your password to access the Administrative Console. The default password is **ecxadLG235**.
5. Click **Manage your certificates**.
6. Select the Certificate Type, browse for the certificate file on your computer, then click **Upload**.
7. Reboot the virtual machine where the application is deployed.

RELATED TOPICS:

- [Microsoft Knowledge Base Article 179380: How to Remove, Import, and Export Digital Certificates](#)
- [Firefox Knowledge Base Article: Advanced settings for accessibility, browsing, system defaults, network, updates, and encryption](#)
- [Google Chrome Knowledge Base Article: Advanced security settings](#)
- [Register a Provider](#) on page **75**

- [LDAP User Name Syntax](#) on page **552**
- [User Administration and Security Management](#) on page **14**

Horizontal Scale Out

In some environments, it may make sense to scale out the **mongod** processes that are part of IBM Spectrum Copy Data Management. Scale out refers to changing the standard deployment, where the two mongod processes are collocated with IBM Spectrum Copy Data Management appliance, to a model where each of the MongoDB database instances run in its own virtual machine (VM).

The general steps to accomplish this are as follows:

- Deploy a MongoDB virtual machine (VM) from the `.ova` file.
- Copy the appropriate VMDKs from IBM Spectrum Copy Data Management to the MongoDB VM.
- Configure the MongoDB VM.

Note: **Steps 1-3**, documented below, will need to be performed twice, once for each of the two MongoDB databases. Proceed to the next step (**4**), only after these steps have been done for each database.

- Configure IBM Spectrum Copy Data Management so that it will use the MongoDB VMs as opposed to the collocated mongod processes running the MongoDB databases.
- Bind directories and edit the `fstab` file to ensure proper function of log collection.

BEFORE YOU BEGIN:

- Ensure IBM Spectrum Copy Data Management is installed and running. See [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page **63**.
- Obtain the VMware `.ova` that deploys a MongoDB VM from Catalogthe `<xxx>`.
- Locate the PowerShell script named `copy_mongodb_vmdk.ps1` that copies the VMDKs from IBM Spectrum Copy Data Management that is included in the `.ova`.
- PowerShell 5.1 or later or PowerCLI 11.2 or later are required to use the `copy_mongodb_vmdk.ps1` script.
- The Node.js script (included in the `.ova`) that manages the MongoDB VM.
- Ensure the account with which you are accessing the vSphere UI has the **Datastore.Browse** privilege.
- Determine two available static IPv4 addresses that can be used for the two VMs that will be deployed.
- If updating from IBM Spectrum Copy Data Management 2.9 to IBM Spectrum Copy Data Management 2.10, the **nfs-utils** module needs to be installed. Run the following command from an IBM Spectrum Copy Data Management appliance terminal window as the **root** user:

```
$ sudo yum --enablerepo=base,updates install -y nfs-utils
```

Step 1: Deploy a MongoDB virtual machine (VM) from the `.ova` file:

1. Log in to the vSphere UI.
2. Deploy the MongoDB VM `.ova` file. **DO NOT** power on the VM.

The resulting VM will have 4 processors and 16GB of RAM with a single 50GB VMDK. The guest operating system is CentOS 7.6.

Step 2: Copy the appropriate VMDKs from IBM Spectrum Copy Data Management to the MongoDB VM.

Note: The next step requires that you know the difference between the two databases. From this point forward, the two databases (and VMs) will be denoted as '**mongo1**' and '**mongo2**'. In a standard deployment case where neither database has been 'extended' (i.e., where neither LVM volume group has had additional logical volumes added to it), mongo1 is the 100GB hard disk, and mongo2 is the 250GB hard disk. In cases where there is not a standard deployment and disks have been extended, knowing which VMware disk belongs to each volume group is critical. It is best to contact support to determine the appropriate VMs.

Each of the steps below will need to be performed for mongo1 and mongo2. Only proceed with mongo2 after completing all of the steps for mongo1.

Prerequisites

- Ensure that both IBM Spectrum Copy Data Management and the MongoDB VM are powered off.
- Have available the IP address of the vCenter server that manages IBM Spectrum Copy Data Management appliance.
- Determine the names of the IBM Spectrum Copy Data Management and MongoDB VM. These are the VM names displayed in the vCenter UI.

1. Open a PowerShell session.
2. In the PowerShell session, navigate (`cd`), into the directory that contains the PowerShell script, `copy_mongodb_vmdk.ps1`.
3. Enter the command, substituting the vCenter IP for **<vCenter_ip>**, the vCenter user name for **<vCenter_userid>**, the IBM Spectrum Copy Data Management VM name for **<IBM Spectrum Copy Data Management_VM_name>**, and the MongoDB VM name for **<MongoDB_VM_name>**:

```
./copy_mongodb_vmdk.ps1 mongo1 <vCenter_ip> <vCenter_userid> <IBM Spectrum Copy Data Management_VM_name> <MongoDB_VM_name> -whatif:$false
```

4. The script will prompt for the hard disk number that corresponds to the database (ex. mongo1). Type the number at the prompt and press **Enter**.
5. When all hard disk numbers have been entered, type the letter '**q**' at the prompt.

Note: Depending on the size of the VMDKs and the speed of the storage network, the script may run for as long as one hour. To observe progress, consult the vSphere UI.

Step 3: Configure the MongoDB VM.

Prerequisites

- A deployed, powered off MongoDB VM to which the IBM Spectrum Copy Data Management VMDK(s) have been copied.
 - Have available the IP address of the IBM Spectrum Copy Data Management virtual appliance.
1. Log in to the vSphere UI.
 2. Power on the MongoDB VM.
 3. After the MongoDB VM is up, use the vSphere UI console facility to open the VM's console.
 4. Log in to the VM's shell using a user name of **ecxadmin** and the password **ecxDP758**. If prompted, set a new password.
 5. Edit (using vi) the file `config.yml` and set the property: `IBM Spectrum Copy Data Management_ip: <IBM Spectrum Copy Data Management_ip>` where **<IBM Spectrum Copy Data Management_ip>** is the IP of the IBM Spectrum Copy Data Management appliance.
 6. In the shell, enter the command: `sudo ./mongoman --func=dsc`. The "Starting MongoDB... done" message will appear.
 7. Enter the command: `sudo ./mongoman`. Based on the VM being configured (mongo1 or mongo2), the "Database Name: mongo1(up)" or "Database Name: mongo2(up)".

Steps 1 - 3 need to be repeated for the second database, mongo2.

Tip: Remember to adjust the database names from mongo1 to mongo2 for the **<MongoDB_VM_name>** variable with executing the `copy_mongodb_vmdk.ps1` PowerShell script and adjust the hard disk number that corresponds to mongo1 to that of mongo2.

Proceed with Step 4 only be after Steps 1 - 3 have been completed for both mongo1 and mongo2.

Step 4: Configure IBM Spectrum Copy Data Management to use the MongoDB VMs rather than the collocated MongoDB processes.

Prerequisites

- Steps 1 - 3 have been completed for both mongo1 and mongo2.
 - A powered off IBM Spectrum Copy Data Management appliance that is in its default configuration with collocated MongoDBs.
 - The password for the IBM Spectrum Copy Data Management appliance's **root** account.
 - Have available the IP addresses of the mongo1 and mongo2 VMs.
1. Log in to the vSphere UI.
 2. Power on the IBM Spectrum Copy Data Management VM through the vSphere UI.

3. After the IBM Spectrum Copy Data Management VM is up, use the vSphere UI console facility to open the VM's console.
4. Log in to the IBM Spectrum Copy Data Management shell as **root** using the password set when you first logged into the IBM Spectrum Copy Data Management shell.
5. Edit (using **vi**) the `managedb.config` file located at `/opt/ECX/tools/scripts/managedb.config` and set the properties:

```
MONGO_HOST=<mongo1_ip>  
MONGO2_HOST=<mongo2_ip>  
MONGO2_PORT=27017
```

where **<mongo1_ip>** is the IP of mongo1 and **<mongo2_ip>** is the IP of mongo2.
6. Restart the IBM Spectrum Copy Data Management appliance. Please give ECX several minutes to fully restart.

Step 5: Bind directories to ensure proper function of log collection.

On the IBM Spectrum Copy Data Management appliance, create and bind the directories that are required for log collection to the newly created mongo1 and mongo2 disks.

1. Log in to the IBM Spectrum Copy Data Management appliance.
2. Open a terminal and enter the following commands, where **<mongo1_ip>** is the IP address of mongo1, and **<mongo2_ip>** is the IP address of mongo2:

```
$ sudo mkdir /tmp/mongo1  
$ sudo mount -t nfs <MONGO1_IP>:/data/log/mongo /tmp/mongo1  
$ sudo mkdir /tmp/mongo2  
$ sudo mount -t nfs <MONGO2_IP>:/data/log/mongo /tmp/mongo2
```

3. To ensure persistence after appliance restart, edit `fstab` to mount the directories automatically. Use **vi** to edit the `fstab` file:

```
$ sudo vi /etc/fstab
```

4. Once the `fstab` file is open, append the following lines:

```
<mongo1_ip>:/data/log/mongo /tmp/mongo1 nfs _netdev 0 0  
<mongo2_ip>:/data/log/mongo /tmp/mongo2 nfs _netdev 0 0
```

5. Save changes to the `fstab` file and exit.

Results

It may take several minutes for IBM Spectrum Copy Data Management to fully restart. Once IBM Spectrum Copy Data Management has restarted, log in to the IBM Spectrum Copy Data Management appliance UI to ensure that the procedure completed successfully.

Restoring A Snapshot From A FlexGroup Volume To Another FlexGroup Volume

Use this process to manually restore a snapshot from a FlexGroup volume to another FlexGroup volume using the OnTap Command Line Interface (CLI).

BEFORE YOU BEGIN:

Consider the following requirements before attempting this procedure.

- Perform the process as a cluster administrator at the **admin** privilege level. If you do not have access to the cluster admin account, using a cluster administrative account with similar access permissions will be necessary.
- The `cifs.show_snapshot` option needs to be set to on. This can be verified from the OnTap CLI on the NetApp controller by issuing:

```
options cifs.show_snapshot
```

- If it is not enabled (off), turn it on by issuing the following command:

```
options cifs.show_snapshot on
```

- For specific FlexGroup requirements, please review the NetApp OnTap Documentation. See the information at <http://docs.netapp.com/ontap-9/index.jsp?topic=%2Fcom.netapp.doc.pow-fg-mgmt%2FGUID-4EABB306-55B2-40C1-93A6-75F538A66833.html>.

To restore a snapshot from a FlexGroup volume to another FlexGroup volume:

If it has not already done so, mount the CIFS share that contains the snapshots to be restored. The `.snapshot` directory will appear at the root of the share. IBM Spectrum Copy Data Management

Note: When viewing contents of the home directory, the directory may be displayed in one of two ways. Likely, if long file names are supported, the directory will be displayed as `~snapshot`, whereas if the operating system does not support long file names, the directory will be displayed as `~SNAPSHT`.

1. Open the OnTap Command Line Interface (CLI).
2. Obtain the name of the source and destination volumes, and the name of the snapshot to be restored.
3. Deactivate any active quota rules on the destination FlexGroup volume. Active quota rules can be deactivated by issuing the following command from the OnTap CLI:

```
volume quota modify -vserver <vserver> -volume <volume> -state off
```

Note: Information contained between the less-than and greater-than symbols should contain the source and destination servers and flexgroups. Omit the less-than and greater-than symbols when running the command. The colon and dashes should remain.

4. Enter the following command to initiate the restore:

```
snapmirror restore -source-path <source vserver>:<source flexgroup> -  
destination-path <destination vserver>:<destination flexgroup> -snapshot  
<snapshot name>
```

5. If there were active quota rules applied to the destination FlexGroup volumes that were deactivated in Step 3, reactivate them. Use the following command to reactivate quote rules:

```
volume quota modify -vserver <vserver> -volume <volume> -state on
```

Documentation and Support

The topics in the following section cover the documentation roadmap, technical support details, and information about the online help system.

Documentation Roadmap

Help System

In IBM Spectrum Copy Data Management, when needed:

- Click the help icon  to invoke help specific to the active function.
- Use the Help system's Search and Index features to locate pertinent information, as these features search the entire documentation suite.

User's Guide

This PDF is intended for IBM Spectrum Copy Data Management users, system administrators, and the Super User. It contains information, procedures, and tips for the most commonly used functions.

System administrators can use this guide to help install, maintain, and start the application, manage users, and catalog resource information. Users can find procedures on how to search and browse for objects, generate and interpret reports, schedule jobs, and orchestrate Backup and Restore jobs.

RELATED TOPICS:

- [About the Help System](#) on page **543**

About the Help System

Starting Help

In the application, click the help icon  to invoke help specific to the active function. For example, click the help icon  on the Search tab to view help topics related to searching for objects.

Pop-up windows must be enabled in your browser to access the Help system and some IBM Spectrum Copy Data Management operations.

Before You Begin, Next Steps, and Related Topics

Prerequisites for procedures are listed in **Before You Begin** sections in many help topics. For example, you must run a job before you run a report, so a link to the Run a Job procedure is available in the Before You Begin section of the Run a Report procedure. Topics also include **Next Steps** and **Related Topics** sections for more information.

Search Help Feature

Use the Search feature in the Help system to locate pertinent information in the entire documentation suite:

- Enter a word in the search field to find all topics that contain that word. For example, schedule.
- Enter multiple words in the search field to find topics related to both words. For example, searching for schedule job returns results for schedule and job as well as "schedule job."
- Enter words separated with AND, +, or & to find topics that contain all of the words separated by the operators. For example, schedule AND job, schedule + job, or schedule & job.
- Enter words separated with OR to find topic that contain any of the words separated by OR. For example, schedule OR job.

Search for words on a help page by using the Find feature in your browser.

Security Management Topics

A security icon  precedes a security management help topic. Security management identifies the interfaces that manage the security functions in IBM Spectrum Copy Data Management. Only the Super User and System Administrators configure the security functions. Examples of security management include adding users, assigning roles, configuring IBM Spectrum Copy Data Management to use LDAP, and configuring IBM Spectrum Copy Data Management to use HTTPS.

RELATED TOPICS:

- [Documentation Roadmap](#) on page 542

Reference Topics

The topics in the following section cover reference information including virtual machine privileges, search and filter guidelines, and return code references.

Virtual Machine Privileges

The following vCenter permissions are required for Backup and Restore operations if a virtual machine is configured as a provider in IBM Spectrum Copy Data Management with credentials lower than an administrator.

vCenter Permissions for Backup and Restore Operations

Datastore

- Allocate space
- Browse datastore
- Low level file operations
- Remove file
- Rename datastore
- Update virtual machine files

Datastore Cluster

- Configure

Folder

- Create folder
- Delete folder

Global

- Cancel task

Host > Configuration

- Storage partition configuration

Network

- Assign network
- Configure
- Move network
- Remove

Profile-driven storage privileges

- Profile-driven storage view

Resource

- Apply recommendation
- Assign a vApp to resource pool
- Assign virtual machine to resource pool
- Migrate powered off VM
- Migrate powered on VM
- Query vMotion

Virtual Machine > Configuration

- Add existing disk
- Add new disk
- Add or remove device
- Advanced
- Change CPU count
- Disk change tracking
- Memory
- Modify device settings
- Raw device
- Reload from path
- Remove disk (detach and remove virtual disk)
- Rename
- Settings

Virtual Machine > Interaction

- Backup operation on virtual machine
- Device connection
- Power Off
- Power On

Virtual Machine > Inventory

- Create from existing
- Register
- Remove
- Unregister

Virtual Machine > Provisioning

- Allow disk access
- Clone virtual machine
- Mark as template
- Mark as virtual machine

Virtual Machine > Snapshot management

- Create snapshot
- Remove snapshot
- Revert to snapshot

vApp

- Add VM to vApp
- Assign resource pool to vApp
- Assign vApp to another vApp
- Create
- Delete
- Power Off
- Power On
- Rename
- Unregister
- vApp resource configuration

In the All Privileges section located below the permissions selection window, ensure **Propagate to children** is selected.

RELATED TOPICS:

- [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page **63**
- [Update IBM Spectrum Copy Data Management](#) on page **524**

Search and Filter Guidelines

Use search and filter fields to tailor the results while conducting a search. The goal of searching and filtering is to provide you with information sets that are manageable and meaningful.

In search and filter fields:

- Enter a character string to find objects with a name that matches or contains the character string. You can also enter partial character strings. Character strings are case insensitive.
- Enter ***** to return all available objects.
- Apply wildcards as needed. Wildcard considerations are described later in this topic.

Perform a basic search using inline search parameters:

Using the following inline search strings, you can perform complex searches based on a file's location, size, and access, creation, or modified time from the basic search field.

Search by object location:

Limit your search to a specific cataloged location using the following examples:

`type:file location:<HOSTNAME>*` searches for all objects on the storage system associated with the entered host

`type:file location:<HOSTNAME>* name:*.txt` searches for .txt files on the storage system associated with the entered host

Search by object size:

Search for cataloged objects with a specific file size or file size range using the following examples:

`size:100KB` searches for all objects that are 100 KB in size

`size:50KB-100MB` searches for all objects between 50 KB and 100 MB in size

`size:*-100MB` searches for all objects that are less than 100 MB in size

`size:100MB-*` searches for all objects that are larger than 100 MB in size

The following size unit strings are supported:

k, K, KB, Kb, kB, kb, KiB, kib, kilobyte, and kilobytes

m, M, MB, Mb, mB, mb, MiB, mib, megabyte, and megabytes

g, G, GB, Gb, gB, gb, GiB, gib, gigabyte, and gigabytes

t, T, TB, Tb, tB, tb, TiB, tib, terabyte, and terabytes

p, P, PB, Pb, pB, pb, PiB, pib, petabyte, and petabytes

Search by object access, creation, and modified time:

Search for cataloged objects that were last accessed, modified, or created at a specific time or time range using the following examples:

`atime:2yearsago` searches for all objects with an access time of two years ago from the time of the search. `ctime` searches against the object's creation time, and `mtime` searches against the object's modification time.

`atime:2yearsago-lastyear` searches for all objects with an access time between last year and two years ago. `ctime` searches against the object's creation time, and `mtime` searches against the object's modification time.

`atime:past2weeks` searches for all objects with an access time from the past two weeks. `ctime` searches against the object's creation time, and `mtime` searches against the object's modification time.

The following time strings are supported:

years, yearsago, year, yearago
months, monthsago, month, monthago
weeks, weeksago, week, weekago
days, daysago, day, dayago
hours, hoursago, hour, hourago
minutes, minutesago, minute, minuteago

Combining search strings:

By combining the above search strings in the basic search field, you can limit your search to specific objects, locations, and size ranges.

```
*.vmdk type:file location:<HOSTNAME>/vmtemplates/* catalog:netapp  
size:2MB-5MB
```

In this example, search results include all resources that include ".vmdk," residing on a resource named <HOSTNAME>/vmtemplates and its subfolders within a NetApp ONTAP catalog, with a size greater than 2 MB but less than 5 MB.

Wildcard Considerations:

A wildcard is a character that you can substitute for zero or more unspecified characters when searching text. Position wildcards at the beginning, middle, or end of a string, and combine them within a string.

- Match a character string with an asterisk, which represents a variable string of zero or more characters:
 - string*** searches for terms like string, strings, or stringency
 - str*ing** searches for terms like string, straying, or straightening
 - *string** searches for terms like string or shoestring

- Match a single character with a question mark:
 - string?** searches for terms like strings, stringy, or string1
 - st??ring** searches for terms like starring or steering
 - ???string** searches for terms like hamstring or bowstring

You can use multiple asterisk wildcards in a single text string, though this might considerably slow down a large search.

RELATED TOPICS:

- [Search for Objects](#) on page **346**
- [View Object Details](#) on page **351**
- [View NetApp ONTAP File Details](#) on page **352**
- [Create an Inventory Job Definition - NetApp ONTAP File](#) on page **186**

Select, Sort, and Reorder Columns

Many tables that open in the user interface are customizable. You can select the columns to display, choose the column to sort on, and revise the order that the columns display.

To select the columns to display:

1. Click the drop-down arrow in the header of any column. Then click **Columns**.
2. Select the columns to display.

To choose the column to sort on:

Click on the header of the column to sort on. If the column is sortable, as indicated by a sort arrow, the table sorts on that column. To sort in the opposite direction, for example ascending versus descending, click the column header again.

To change the order that the columns display:

Drag the header of any column left or right to the location where you want it to appear.

RELATED TOPICS:

- [View a Provider](#) on page **95**
- [Edit a Schedule](#) on page **157**
- [Edit a Job Definition](#) on page **342**
- [Monitor a Job Session](#) on page **164**
- [Search Overview](#) on page **345**
- [Run a Report](#) on page **361**

LDAP User Name Syntax

LDAP form is used when setting up LDAP authentication in IBM Spectrum Copy Data Management. The following shorthand is used for LDAP form:

- cn: Canonical Name
- dn: Distinguished Name
- rdn: Relative Distinguished Name
- ou: Organizational Unit
- dc: Domain Component

Following is a typical username entry in LDAP form:

cn=administrator, cn=users, dc=company, dc=com

RELATED TOPICS:

- [Upload an SSL Certificate](#) on page **532**
- [User Administration and Security Management](#) on page **14**

Return Code Reference

Return Codes are issued when a script generated by a user-defined script is run. As the script runs, IBM Spectrum Copy Data Management interprets a return code of 0 as success and return codes 128-255 indicate that the command abnormally terminated. The formula 128+N is used with N representing the signal the process terminated on, for example, return code 143 indicates signal number 15 is caught and causes the executable to terminate abnormally.

Return Code Values

Return Code (128+N)	Return Code Value
1	SIGHUP
2	SIGINT
3	SIGQUIT
4	SIGILL
5	SIGTRAP
6	SIGABRT
7	SIGBUS
8	SIGFPE
9	SIGKILL
10	SIGUSR1
11	SIGSEGV
12	SIGUSR2
13	SIGPIPE
14	SIGALRM
15	SIGTERM
16	SIGSTKFLT
17	SIGCHLD
18	SIGCONT
19	SIGSTOP

Return Code (128+N)	Return Code Value
20	SIGTSTP
21	SIGTTIN
22	SIGTTOU
23	SIGURG
24	SIGXCPU
25	SIGXFSZ
26	SIGVTALRM
27	SIGPROF
28	SIGWINCH
29	SIGIO
30	SIGPWR
31	SIGSYS
34	SIGRTMIN
35	SIGRTMIN+1
36	SIGRTMIN+2
37	SIGRTMIN+3
38	SIGRTMIN+4
39	SIGRTMIN+5
40	SIGRTMIN+6
41	SIGRTMIN+7
42	SIGRTMIN+8
43	SIGRTMIN+9
44	SIGRTMIN+10
45	SIGRTMIN+11
46	SIGRTMIN+12
47	SIGRTMIN+13

Return Code (128+N)	Return Code Value
48	SIGRTMIN+14
49	SIGRTMIN+15
50	SIGRTMAX-14
51	SIGRTMAX-13
52	SIGRTMAX-12
53	SIGRTMAX-11
54	SIGRTMAX-10
55	SIGRTMAX-9
56	SIGRTMAX-8
57	SIGRTMAX-7
58	SIGRTMAX-6
59	SIGRTMAX-5
60	SIGRTMAX-4
61	SIGRTMAX-3
62	SIGRTMAX-2
63	SIGRTMAX-1
64	SIGRTMAX

RELATED TOPICS:

- [Create a Script Job Definition](#) on page 339

Frequently Asked Questions

The following are answers to frequently asked questions related to IBM Spectrum Copy Data Management functionality. The questions and answers are organized by deployment, resources, connectivity, cataloging, operation, and control topics.

- Review the frequently asked questions by topic:
 - For questions related to IBM Spectrum Copy Data Management installation and logging in, see [Deployment](#) on page **556**.
 - For questions related to resources used by IBM Spectrum Copy Data Management, see [Resources](#) on page **557**.
 - For questions related to connecting to IBM Spectrum Copy Data Management resources, see [Connectivity](#) on page **558**.
 - For questions related to IBM Spectrum Copy Data Management cataloging and jobs, see [Cataloging](#) on page **559**.
 - For questions related to general IBM Spectrum Copy Data Management operation, see [Operation](#) on page **560**.
 - For questions related to Backup and Restore jobs in IBM Spectrum Copy Data Management, see [Backup/Restore Jobs](#) on page **561**.

Deployment

How is IBM Spectrum Copy Data Management distributed?

In most cases, IBM Spectrum Copy Data Management is distributed as a virtual appliance through an OVF template.

How do I configure IBM Spectrum Copy Data Management out-of-box?

See the topics [Deployment Checklist](#) on page **19** and [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page **63**.

What are the requirements of the datastores used for the hard disks? What types of VMware datastores are supported?

The type of datastore on which IBM Spectrum Copy Data Management is deployed is transparent to IBM Spectrum Copy Data Management.

Can the virtual appliance hard disks be thin provisioned?

Yes. This is a function of the virtual appliance, and can be set during IBM Spectrum Copy Data Management installation. Better performance can be achieved with thick provisioning of the appliance.

Why is there a delay from when the machine boots to login?

When you boot the machine, several processes occur including:

- Operating system and network connections initiate.
- Dependencies are scanned and resolved.

The more heavily loaded the ESX server is, the longer the boot process might take.

What are the default IBM Spectrum Copy Data Management user names and passwords?

When logging on to IBM Spectrum Copy Data Management for the first time, the default user name is **admin** and the default password is **password**. You will be prompted to reset the default password.

When logging on to the management console of the virtual machine, the default user name is **administrator** and the default password is **ecxadLG235**.

How is the *root* password secured?

You are prompted to change the *root* password on the first *root* login.

Resources

Can the disks be increased dynamically?

IBM Spectrum Copy Data Management data volumes can be expanded if necessary with the approval of Technical Support.

What resource can I add to improve IBM Spectrum Copy Data Management performance?

Increasing memory should help improve performance.

How much of the virtual machine default configuration can be modified?

Parameters such as network, CPU, and memory can be configured at the virtual machine level, but adjusting to below the default levels may impact performance. For default requirements, see the topics [System Requirements](#) on page 23 and [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page 63.

Is it possible to install proprietary software, such as antivirus software, on the virtual appliance?

It is not recommended to install third party applications on the virtual appliance without approval from Technical Support.

Can I access the IBM Spectrum Copy Data Management user interface remotely?

Yes. The IBM Spectrum Copy Data Management user interface is browser based. Supported browsers and the URL are described in the topics [System Requirements](#) on page 23 and [Start IBM Spectrum Copy Data Management](#) on page 65.

What ports are needed to access the IBM Spectrum Copy Data Management user interface?

To access IBM Spectrum Copy Data Management, appropriate ports need to be opened through the firewall. For details, see the topic [User Administration and Security Management](#) on page 14.

What operating system is IBM Spectrum Copy Data Management built on?

CentOS is the operating system on the IBM Spectrum Copy Data Management virtual appliance.

Is Java used as part of the IBM Spectrum Copy Data Management appliance?

Yes. However, OpenJDK is used as opposed to JRE.

Is JavaScript required for accessing IBM Spectrum Copy Data Management?

A browser that supports JavaScript is required. Supported browsers and the URL are described in the topics [System Requirements](#) on page 23.

Do the IBM Spectrum Copy Data Management cataloging and reporting functions impact the performance of the registered storage systems?

The cataloging function is built on technology that is designed to run as low priority on the storage system and automatically adjust itself to give top priority to primary workload operations. The reporting functions do not impact registered storage systems as they run on the IBM Spectrum Copy Data Management virtual appliance.

Connectivity**How does IBM Spectrum Copy Data Management connect to NetApp ONTAP storage systems?**

IBM Spectrum Copy Data Management connects to NetApp ONTAP storage systems through HTTPS or HTTP.

Is the network traffic secure?

Network traffic between the IBM Spectrum Copy Data Management virtual appliance and the the IBM Spectrum Copy Data Management user interface is secured using HTTPS protocol. Network traffic between the IBM Spectrum Copy Data Management virtual appliance and an external resource, such as a NetApp storage system, a vCenter, or an LDAP server, uses either HTTPS or HTTP protocol, which is decided by the System Administrator when registering the resource.

Does IBM Spectrum Copy Data Management work with storage vendors other than DellEMC, IBM, and NetApp?

IBM Spectrum Copy Data Management software works with DellEMC, IBM, and NetApp storage and VMware infrastructure.

IBM Spectrum Copy Data Management provides Backup and Restore support for customers with VMware leveraging heterogeneous storage, extending Backup use cases to VMware on mixed storage.

For the search and reporting features of IBM Spectrum Copy Data Management, the VMware environment can use any storage; it does not have to be DellEMC Unity, IBM, or NetApp. Therefore, IBM Spectrum Copy Data Management provides visibility and insight into VM information across any storage device.

Does IBM Spectrum Copy Data Management work with volumes that have non-Windows file systems?

Yes. IBM Spectrum Copy Data Management catalogs NFS and CIFS files residing on NetApp ONTAP volume snapshots. Linux/Unix files are stored using NFS and Windows files are stored using CIFS protocol.

Does IBM Spectrum Copy Data Management work with SnapManager data?

IBM Spectrum Copy Data Management catalogs the meta-data on LUNs created by SnapManager for SQL Server and Exchange. File level granularity of content hosted inside these LUNs is not available.

Is IBM Spectrum Copy Data Management software SNMP compliant?

Not at this time.

Cataloging

If you add a vCenter into an Inventory job definition, does that automatically discover all the ESX servers within that vCenter?

Yes. Once cataloged, view available VMware resources through the Inventory browse function on the Search tab.

How many storage systems can be cataloged?

IBM Spectrum Copy Data Management can catalog any number of storage systems and is only limited by the data disk in its delivered configuration.

When a catalog job runs, is it a full catalog job each time?

Yes. A full catalog job, not an incremental, is run each time.

Why is it that sometimes many jobs and tasks are marked with Waiting indicators on the Jobs tab?

The number of operations in progress on the Jobs tab varies depending on the number of jobs currently running. IBM Spectrum Copy Data Management controls the number of jobs allowed to run. When the number of jobs exceeds the value defined by IBM Spectrum Copy Data Management, jobs marked with Waiting indicators display on the Jobs tab. IBM Spectrum Copy Data Management also controls the number of job tasks to run simultaneously for a given job when multiple jobs are running.

When does the cataloged data get cleaned up?

After a certain number of runs for a given job, older objects for that job are purged from the Catalog. This retention parameter is set when the job is defined.

The Maintenance job removes resources and associated objects created by IBM Spectrum Copy Data Management when a job in a pending state is deleted. The cleanup procedure reclaims space on your storage devices, cleans up your IBM Spectrum Copy Data Management catalog, and removes related snapshots.

See the topic [Maintenance Job](#) on page 336.

Operation

How do I protect and recover the IBM Spectrum Copy Data Management appliance itself?

Backing up the IBM Spectrum Copy Data Management appliance regularly is a critical operation. For more information, contact Technical Support.

Can I restore an individual file by using the Search window?

Yes. To restore a file on a Windows machine, search for the file in IBM Spectrum Copy Data Management and discover its location through the file's properties pane. See the topic [Find and Restore a File](#) on page 354.

The recommended Best Practice for protection/recovery is to use Backup and Restore jobs in IBM Spectrum Copy Data Management.

Why is it that when I select Hide Duplicates when performing an advanced search, some duplicate objects still display in the search results pane?

In some cases, the name of a returned object on the search results pane may be the same as another object, however the resources where the objects reside is different. Review the file properties of the objects by selecting their names on the search results pane to view the differences between the returned entries.

Can I generate customized reports?

IBM Spectrum Copy Data Management provides a set of predefined reports that can be customized through the use of parameter selection.

How are IBM Spectrum Copy Data Management logs collected for troubleshooting?

There are two approaches for downloading logs. Download logs from the Support menu or access the IBM Spectrum Copy Data Management appliance through a command prompt. The first approach is simpler and generally sufficient. The second approach produces a more comprehensive set of logs. See the topic [Collect Logs For Troubleshooting](#) on page 518.

Is audit tracking provided?

An audit log displaying IBM Spectrum Copy Data Management activity is available through the **Support** menu. Click the arrow next to the **Support**  icon, then click **View Audit Log**.

Backup/Restore Jobs

For a Backup job, how do I update the retention after a job has run?

Open the existing job definition, click Snapshot in the workflow pane, and update the Keep Snapshots parameter. The retention policy changes to the supplied value when the job is next run.

To what extent do the IBM Spectrum Copy Data Management Backup and Restore functions impact the performance of NetApp storage systems?

The Backup and Restore functions employ technologies such as Snapshot and FlexClone that are designed to be low-impact on the NetApp storage systems. Generally, users should observe little unexpected performance impact on the storage systems.

For a Backup job, can I exclude swap partitions?

No. The lowest granularity of protection is a virtual machine.

RELATED TOPICS:

- [Oracle Database Support FAQ](#) on page **563**
- [Microsoft SQL Server Support FAQ](#) on page **584**
- [System Requirements](#) on page **23**

Oracle Database Support FAQ

What is Oracle CDM? How does it help solve my challenges?

Database administrators working extensively with Oracle are challenged when faced with mission-critical use cases such as Backup, Recovery, DevOps, and Business Analytics. This is especially true given that their Oracle databases have expanded in size and number over time, and that the databases need to be up and running 24x7x365.

Oracle DBAs struggle with the following:

- Backups are slow, complex and need constant management
- Backup process slows down the production servers
- Recoveries are slow and complex
- Repurposing App consistent backups (clones) for DevOps and Business Analytics is slow, complex and storage inefficient
- Lack of automation exists for producing quick and secure clones required to accelerate DevOps
- Copy sprawl problems occur due to no central catalog of copies
- Unable to meet organization stringent RPO and RTO requirements

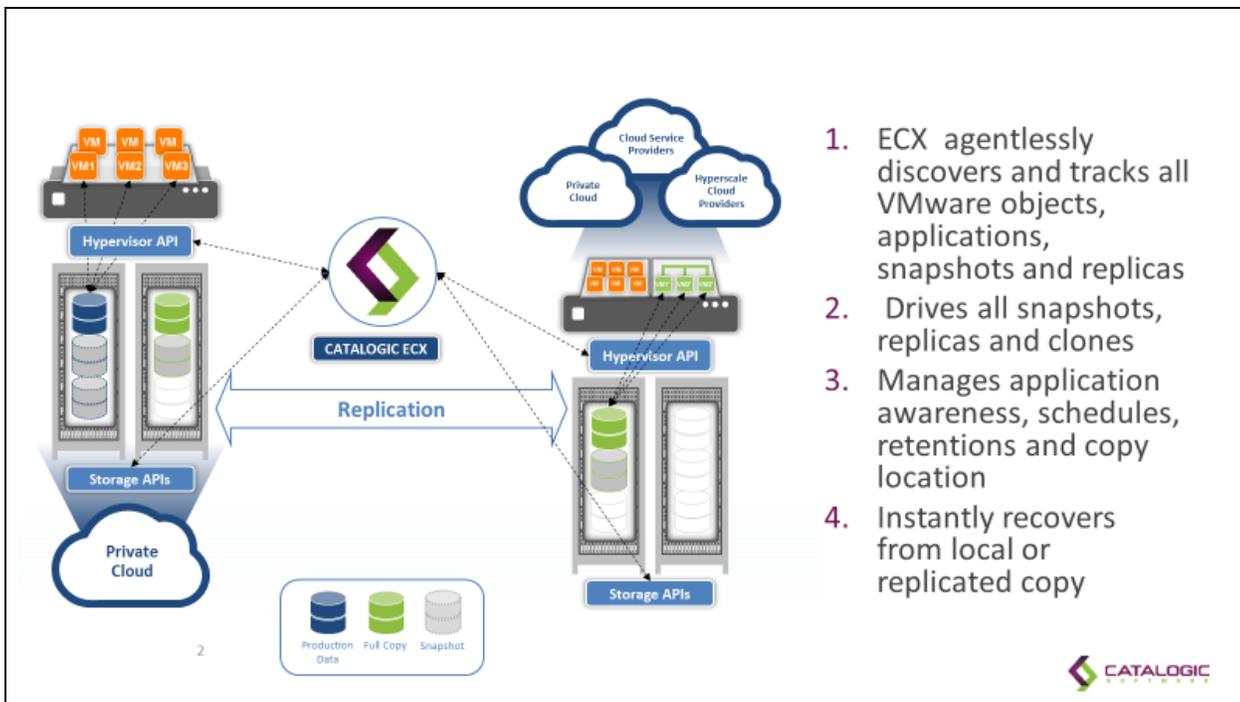
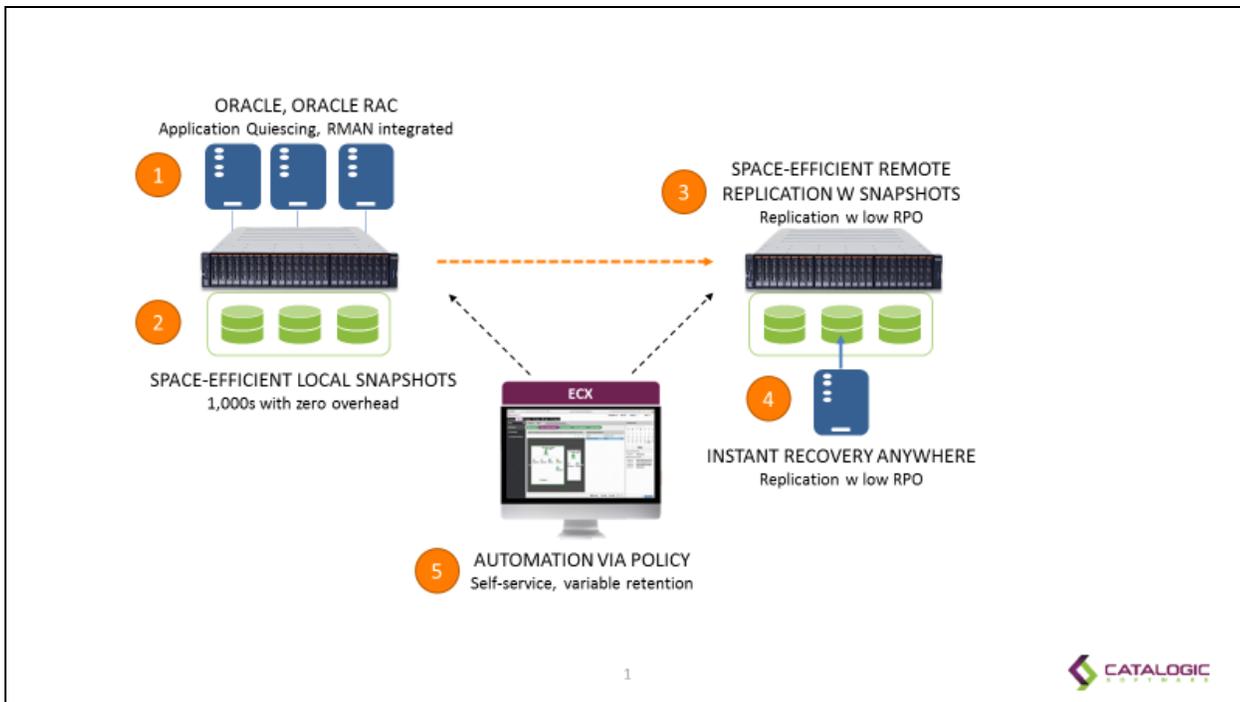
IBM Spectrum Copy Data Management simplifies Oracle database copy management by enabling administrators to orchestrate application-consistent copy creation, cloning and recovery in minutes, instead of hours or days. IBM Spectrum Copy Data Management copy management leverages the advanced snapshot and replication features of the underlying storage platform to rapidly create, replicate, clone, and restore copies of Oracle databases in the most efficient way possible, in both time and space. IBM Spectrum Copy Data Management enables you to focus on the backup and recovery requirements of your business rather than the technical details of the underlying storage platforms.

IBM Spectrum Copy Data Management is an intelligent copy data management solution that delivers end-to-end automation, orchestration, and self-service functionality for your Oracle environment through a comprehensive and scalable Inventory. With the self-service features of IBM Spectrum Copy Data Management, your users are empowered to create clones on demand, freeing DBAs, while at the same time offering the advanced recovery features needed for Oracle environments.

IBM Spectrum Copy Data Management Oracle Copy Data Management solution supports the following Oracle deployment modes:

- Single instance – a single instance running on a single server accessing a database
- RAC (Real Application Clusters) leveraging ASM – more than one instance running on multiple servers are accessing a database simultaneously
- ASM (Automatic Storage Management) – Oracle's own volume manager and cluster filesystem that is optimized for Oracle Database features.

IBM Spectrum Copy Data Management Solution Architecture



1. ECK agentlessly discovers and tracks all VMware objects, applications, snapshots and replicas
2. Drives all snapshots, replicas and clones
3. Manages application awareness, schedules, retentions and copy location
4. Instantly recovers from local or replicated copy

Deployment and Registration

Do I need to deploy any additional agents to protect Oracle standalone or Oracle RAC servers?

IBM Spectrum Copy Data Management for Oracle is delivered as a VMware OVA that is easily deployed on demand in a matter of minutes. Once deployed, you simply register your Oracle servers with appropriate credentials and then let IBM Spectrum Copy Data Management discover the rest. IBM Spectrum Copy Data Management eliminates the complexity of manually deploying and maintaining application agents on Oracle servers. A lightweight application-aware component is automatically injected to the required Oracle servers on demand and automatically updated to the latest version if required.

Oracle Backup creation workflow

App consistent Oracle database backup creation (local and remote) step by step

IBM Spectrum Copy Data Management auto-discovers databases and enables copies only of eligible databases. To be eligible for IBM Spectrum Copy Data Management backup, the Oracle database needs to be residing on a supported storage platform. With IBM Spectrum Copy Data Management, application owners do not need to be concerned about storage infrastructure.

IBM Spectrum Copy Data Management creates application-consistent Oracle database copies without the need to build and maintain complex RMAN scripts.

A typical IBM Spectrum Copy Data Management Oracle database backup creation workflow consists of the following steps:

- Auto-inject a lightweight component into the standalone Oracle Server(s) or one of the Oracle RAC server node(s) running a database instance to be copied
- Discover storage volume mapping to selected Oracle database(s) and logs
- Place the Oracle database in hot backup mode
- Automatically create a consistency group for related storage volumes
- Create an application-consistent backup of the consistency group
- Take the Oracle database out of hot backup mode (typically within a few seconds of entering the hot backup mode)
- Optionally create log copies into the specified mount points
- Optionally create selected masked copies using masking tools for secure DevOps use
- Catalog Oracle copies in Inventory and optionally record details in RMAN recovery catalog
- Optionally replace application-consistent backup to remote location leveraging storage replication
- Clean up auto-injected components from Oracle server node

IBM Spectrum Copy Data Management creates and uses in-place copies, so no data is physically moved. Replication to off-host storage is performed by leveraging storage replication, which reduces the amount of impact on Oracle Servers and Databases. IBM Spectrum Copy Data Management generated application-consistent copies are both space and time efficient. With the same ease, a DBA can automate the creation of remote copies for DR use cases.

Does IBM Spectrum Copy Data Management Oracle solution leverage the Storage Consistency Group feature?

The storage consistency group feature allows storage administrators to take a snapshot of database applications where the data is spread across multiple volumes to maintain consistency across all volumes.

In a typical Oracle Database, the data is spread across different volumes for better IO performance and availability. IBM Spectrum Copy Data Management Oracle application-consistent backup creation ensures that appropriate consistency groups are automatically created to maintain consistency across all related volumes.

What level of selection granularity is supported for Oracle backup workflows?

IBM Spectrum Copy Data Management Oracle backup workflows support copy selection at the following levels:

- One or more Oracle home locations
- One or more databases
- One or more container databases for Oracle 12c

Why can I not select some of the databases for protection in an Oracle backup workflow?

You cannot select a database if it is not eligible for protection. Hover your cursor over the database name to view the reasons the database is ineligible, such as that the database files, control files, or redo log files are stored on unsupported storage.

Will IBM Spectrum Copy Data Management auto discover newly added databases and automatically protect them?

If you select the parent Oracle Home in a Backup job definition, all databases under it are protected. If a new database is added under the home, it will be automatically protected once it is cataloged. Discovery and cataloging of new Oracle databases occurs as part of regularly scheduled Oracle Inventory job.

Does the Oracle database need to be on supported storage for IBM Spectrum Copy Data Management?

Yes, the IBM Spectrum Copy Data Management Oracle solution leverages storage snapshots for database protection. All databases, database files, control files and redo logs must be on supported storage systems for it to be eligible for protection.

Does IBM Spectrum Copy Data Management leverage the Oracle 12c Storage Snapshot feature?

This new feature of Oracle 12c enables you to take a storage snapshot of your database without needing the database to enter BACKUP mode. In Oracle, when you need to recover, you can use a point in time of the snapshot. You can roll forward by using the database archive logs, and use this snapshot feature to recover part or all of the database. IBM Spectrum Copy Data Management fully supports this feature starting in IBM Spectrum Copy Data Management 2.5.1.

Does IBM Spectrum Copy Data Management support protection of Offline Databases?

Databases in offline mode are not automatically included during backup workflows, unless they share storage volumes with a selected database that is active. IBM Spectrum Copy Data Management marks the offline databases and does not present them for Oracle Restore Workflows. Users may be able to retrieve these database files as flat files and perform application mount outside of IBM Spectrum Copy Data Management.

Does IBM Spectrum Copy Data Management support protection of Oracle databases not running in Archive Log mode?

Yes, protection of Oracle Databases running in NOARCHIVELOG mode are now supported for both Inventory and Backup use cases. You can recover database to the point of the most recent snapshot. PIT recoveries are not supported for databases running in NOARCHIVELOG mode.

Does IBM Spectrum Copy Data Management support protection of Oracle databases using pFile (text initialization parameter files)?

Yes, IBM Spectrum Copy Data Management now supports IBM Spectrum Copy Data Management backup of databases started through pFile in addition to spFile.

Oracle Archive log management**Does IBM Spectrum Copy Data Management support archive log backup and log management?**

Oracle DBMS creates database transaction logs as part of its operation. Oracle databases can run in the following logging modes:

- NOARCHIVELOG mode – In NOARCHIVELOG mode, no transaction logs are created, and there is no capacity to run point-in-time recovery or online backups. This is the default.
- ARCHIVELOG mode – In ARCHIVELOG mode, the database makes copies of all online redo logs after they are filled. These copies are called archived redo logs. The archived redo logs are created via the ARCH process. The ARCH process copies the archived redo log files to one or more archive log destination directories. These saved archived logs are used for point-in-time recovery.

IBM Spectrum Copy Data Management provides you with an option for archive log files processing:

- Enable archive log backup (Recommended)
- Use existing archive log (Default)

IBM Spectrum Copy Data Management greatly simplifies archive log protection. If a user chooses to protect archive logs, IBM Spectrum Copy Data Management enables continuous backup of archive logs to a specified destination providing the lowest RPO (transaction level recoveries).

IBM Spectrum Copy Data Management automatically discovers the location where Oracle writes archived logs. If this location resides on storage from a supported vendor, IBM Spectrum Copy Data Management can protect it. If the existing location is not on supported storage, or if you wish to create an additional backup of database logs, enable the Create Additional Archive Log Destination option in the Oracle Backup job definition, then specify a path that resides on supported storage. When enabled, IBM Spectrum Copy Data Management configures the database to start writing archived logs to this new location in addition to any existing locations where the database is already writing logs. If multiple databases are selected for backup, then each of the servers hosting the database must have their destination directories set individually.

Can I specify a retention period for backed up archive logs within IBM Spectrum Copy Data Management?

If the Create Additional Archive Log Destination option is selected, IBM Spectrum Copy Data Management automatically manages the retention of only those archived logs that are under the new destination specified in the job definition. After a successful backup, logs older than the backup are automatically deleted from the IBM Spectrum Copy Data Management-managed destination.

If the Use Existing Archive Log Destination(s) option is selected in Oracle Backup job definition, IBM Spectrum Copy Data Management does not automatically purge any archived logs. The retention of archived logs must be managed externally, for example using RMAN. In order to support point-in-time recovery, ensure that the retention period is at least large enough to retain all archived logs between successive runs of the Oracle Backup job.

Oracle RMAN integration

Does the IBM Spectrum Copy Data Management Oracle solution integrate with RMAN?

Oracle Recovery Manager (RMAN), a command-line and Enterprise Manager-based tool, is the method preferred by Oracle DBAs for backup and recovery of Oracle databases, including maintaining an RMAN repository.

IBM Spectrum Copy Data Management creates application-consistent Oracle database copies simply – with no need to build and maintain complex RMAN protection scripts. At the same time, IBM Spectrum Copy Data Management automates cataloging of Oracle database copies in the RMAN recovery catalog. This enables DBAs to leverage RMAN for:

- Verification - IBM Spectrum Copy Data Management-created Oracle database copies can be instantly mounted so they can be easily verified through the RMAN verify command
- Advanced Recovery - IBM Spectrum Copy Data Management-created Oracle database copies can be instantly mounted to perform RMAN-driven PIT recoveries of database and tablespace.

IBM Spectrum Copy Data Management offers the following choice to the user for RMAN catalog registration:

- Register all copies in the RMAN catalog to enable RMAN recoveries against all copies
- Register on-demand only selected copies in the RMAN catalog to enable RMAN recoveries only when you need to recover

IBM Spectrum Copy Data Management simplifies full application-aware Oracle-consistent copy lifecycle while maintaining the flexibility and benefits of full RMAN-driven advanced recovery capabilities.

Pre and Post Scripts

Does IBM Spectrum Copy Data Management support pre/post scripts for the Application Backup workflow?

Yes, IBM Spectrum Copy Data Management supports job-level Pre/Post scripts and job-level pre/post Snapshot scripts to enable further customization.

- Job-level prescripts and postscripts are scripts that can be run before or after a job runs.
- Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot subpolicy runs. (Please refer to pre/post script topic in the IBM Spectrum Copy Data Management User's Guide for details.)

Data Masking

Does the IBM Spectrum Copy Data Management Oracle solution offer Data Masking integration with third party masking tools?

A concern for security officers in any organization is to keep confidential information locked down, even internally. Data masking is used to hide confidential data, by replacing it with fictitious data, when making data copies for DevTest or other use cases. It prevents leakage of sensitive data in non-production databases via static data masking [SDM], and production data in transit via dynamic data masking [DDM].

IBM Spectrum Copy Data Management includes integrated data masking workflows with the ability to leverage third party masking tools. Traditionally, data masking is difficult, slow, and storage-consuming, but with IBM Spectrum Copy Data Management it is easily integrated into the Oracle Backup workflow, allowing creation of masked copies at a specified frequency. Masked copies are automatically marked in the Inventory. Access to secure copies is managed by the administrator by leveraging the application-level RBAC.

Is sample masking script provided with IBM Spectrum Copy Data Management?

A sample data masking script can be provided upon request. A sample masking script demonstrates data masking integration with built-in data masking functionality of Oracle 11g and 12c database system.

Oracle Restore Workflow

Can I leverage Oracle database clones for multiple use?

Limitations of current tools and approaches:

- Database cloning requires action by DBAs and is gated by process
 - QA relies on DBAs for cloning the databases for functional testing
 - The database cloning is gated by processes (space requisition, approvals, etc.)
- Database cloning is not time- or storage-efficient
 - Usage of RMAN or custom scripts creates full backup requiring large amounts of additional storage
 - Creating full copies is slow

IBM Spectrum Copy Data Management solves these challenges with a simple, automated, end-to-end clone lifecycle management:

- Self-service access to secure clones by QA team eliminates administrative and process bottlenecks
- IBM Spectrum Copy Data Management enables rapid database clones that are both time- and space-efficient
 - Provisions clones in minutes regardless of their size
 - Leverages underlying storage snapshots for space efficiency
- IBM Spectrum Copy Data Management promotes standardization and governance through centralized Inventory, granular RBAC, and automated jobs

Your Oracle clones can be utilized and consumed instantly – for whatever your use case -- through IBM Spectrum Copy Data Management “Instant Disk Restore” jobs. IBM Spectrum Copy Data Management catalogs and tracks all cloned instances. Instant Disk Restore can leverage iSCSI or FC protocol to provide an immediate mount of LUNs without transferring data.

Can I create an Instant Clone of an Oracle database for DevOps and Business Analytics?

IBM Spectrum Copy Data Management provides automated workflows to create instant clones of Oracle database regardless of its size.

- Instantly create database clones from any of the copies in the IBM Spectrum Copy Data Management Inventory, at local or remote locations, to accelerate Business Analytics.
- Enable and accelerate DevOps by providing Instant Disk Restore to secure clones of databases to appropriate users via application-level RBAC.

Then, when your TestDev, DevOps, or research/analytics work is completed, you can save the clone to more permanent storage or simply tear it down.

Can I create multiple clones from the same database snapshot?

Yes, multiple copies can be provisioned by IBM Spectrum Copy Data Management from the same snapshot. These copies can be either mounted with new DB names on the same Oracle Server or with same/different names on different Oracle servers.

How much space is used by my Oracle clones?

IBM Spectrum Copy Data Management utilizes zero-footprint instant clone capabilities provided by storage arrays. These clones do not use any additional space (except for maintaining pointers). Only new changes/writes to the read/writable clones utilize space.

What is the granularity of Database recovery supported by IBM Spectrum Copy Data Management?

Supported recoveries for standalone or RAC configurations

- Instant recovery of Oracle databases regardless of the size of the database
- Recover database(s) to original or to a new server (physical or virtual) simply with few clicks

- Recover database(s) with new names simply with few clicks
- Recover at DR site from replicated copies simply with few clicks
- Database can be recovered to point of snapshots to original or new location
- Perform PIT recoveries to original or new location simply with few clicks
- Recover RAC databases to any node on RAC simply with few clicks
- Database running on older versions can be recovered to an instance running same or newer version (IBM Spectrum Copy Data Management inherits any limitations defined by Oracle)
- Each selected database in a Restore job can have a separate destination specification. Databases can be recovered using a new name to an original or new instance
- You can select one or more databases in a single Restore job definition
- Recoveries are supported across the same storage type (i.e. ASM to ASM and standalone to standalone)
- Databases are always recovered in online mode

Users can additionally perform advanced fine-grained recoveries via RMAN integration. You can use IBM Spectrum Copy Data Management to instantly mount required snapshot copies to a specified Oracle server and use RMAN to recover. All RMAN-supported recoveries are available to users.

What granularity of Point-in-Time recovery is supported?

IBM Spectrum Copy Data Management enables you to recover databases to a specific point-in-time to enable you to:

- Restore to the state just before the point of failure.
- Restore multiple databases to a consistent time.

During PIT (point-in-time) Oracle database recoveries, if no log file snapshot exists that is newer than the chosen recovery PIT, the job creates a fresh log snapshot and uses it for recovery.

What happens to Database SID during recovery?

The Oracle System ID ([SID](#)) is used to uniquely identify a particular database on a server. One cannot have more than one database with the same SID on a single server. When using [RAC](#), all instances belonging to the same database must have unique SIDs. A Restore job definition provides the user with an option to define a new name for a database during recovery. IBM Spectrum Copy Data Management will rename the SID if the user chooses a new database name during recovery, otherwise the original SID will be used.

How is the Oracle Initialization parameter file (init.ora) processed during Database recovery?

The Oracle initialization parameter file (init.ora) is created by the DBA and defines the overall instance configuration, such as how much memory should be allocated to the instance, the file locations, and internal optimization parameters.

IBM Spectrum Copy Data Management catalogs the initialization parameters during a job and uses them during recovery. IBM Spectrum Copy Data Management provides advanced options to control the behavior of

processing initialization parameters used to start up the recovered database in Instant Database Recovery and DevOps workflows.

Customizable options allow you to use the same parameters as the source or specify a template pfile file to use.

Additionally, cluster-related parameters like `instance_number`, `thread`, and `cluster_database` are set automatically by IBM Spectrum Copy Data Management depending on the appropriate values for the destination.

For more information, see [Restore Jobs - Rename Mount Points and Initialization Parameter Options](#) on page **331**.

How does IBM Spectrum Copy Data Management set up mount points during Database recovery?

IBM Spectrum Copy Data Management provides advanced options for Instant Database Recovery and DevOps workflows to override the behavior of mount point handling during recovery. The Mount Point Rename option provides these selections:

- Append a timestamp: IBM Spectrum Copy Data Management appends a timestamp to the original mount point.
- Do not rename: IBM Spectrum Copy Data Management uses the same path/name for mount points or ASM diskgroups as the source.
- Add a custom prefix: Select this option to specify a custom prefix to be prepended to the source paths/names. The prefix value may contain leading or trailing slashes. In the case of ASM diskgroup names, the slashes are removed.
- Add a custom suffix: Select this option and specify a custom suffix to be appended the source paths/names.
- Replace a substring: Select this option to specify a custom string of characters in the old mount point to be replaced with another string of characters.

For examples of these options, see [Restore Jobs - Rename Mount Points and Initialization Parameter Options](#) on page **331**.

Can I recover an Oracle Database running on a Linux physical server to Oracle running on a Linux VM?

Yes, the recovered database will be recovered as Physical RDM (pRDM).

Can I recover an Oracle Database running in Linux VM to Oracle running on a Linux Physical Server?

Yes, the pRDM configured database will be recovered as LUNs on the Physical Linux Oracle server. Oracle configured with VMDK can only be recovered to VM.

Can I set up an Oracle Database to automatically be refreshed on a schedule?

The database "refresh" can be achieved by checking the "Allow overwrite and force cleanup of old session" option in the Restore job definition. You can run a job that spins up a database and the session will go into a

Pending state awaiting cleanup. You can then run the same job again at a later time (either manually or scheduled) and it will automatically clean up the previous session before starting the new one. The cleanup process automatically stops and dismounts the database from the previous session. The new session will then mount and start the database again from the chosen snapshot (which is typically the latest as of runtime).

Where are Oracle specific and IBM Spectrum Copy Data Management specific logs if errors occur?

All required logs (IBM Spectrum Copy Data Management and Oracle application) are collected as part of the current log collection functionality. There should be no need to manually obtain Oracle application logs from within Oracle Servers.

Can a persistent agent be installed?

All required plugins are automatically injected on demand and automatically updated to the latest version if required.

Does Oracle application level encryption, such as Transparent Data Encryption(TDE), impact IBM Spectrum Copy Data Management?

Transparent Data Encryption (TDE) stops would-be attackers from bypassing the database and reading sensitive information from storage by enforcing data-at-rest encryption in the database layer. This is application-layer encryption that wouldn't typically impact IBM Spectrum Copy Data Management. The encryption keys live outside the database in a "wallet" that is managed separately by the administrator. IBM Spectrum Copy Data Management will create copies of the data on server A and mount them on server B, for example. The database software on server B should be able to read the encrypted data as long as the necessary keys are installed in the wallet there.

How do I refresh? How do I promote to Production?

All database recovery operations can leverage Instant Recovery mode (Test) and can then either be deleted or promoted to permanent mode via workflow control. This behavior is controlled via the job definition option Make Permanent.

- Enabled - Always make permanent
- Disabled - Never make permanent
- User election - Allows the user to select Make Permanent or Cleanup when the job session is pending

Does IBM Spectrum Copy Data Management provide Oracle specific reporting?

Reports are offered to assure that your Oracle databases are sound and that your IBM Spectrum Copy Data Management jobs are verified. Reports provided by IBM Spectrum Copy Data Management specifically for application support include:

- Application Configuration Report, which describes valuable system information about your Oracle Database Servers, and affirms that Oracle is configured correctly to be eligible for backup creation.
- Application RPO Compliance Report, which determines which of your Oracle database servers are not in compliance with your RPO parameters, and displays the reasons for their non-compliance.

System Requirements

What Oracle versions are supported and on what OS?

Oracle Configuration Requirements

Database Versions/Types	Server Types	Operating Systems	Storage Configuration	Storage Systems
<p>Oracle 11g R2, 12c R2, or 18c [1] [13] configured as:</p> <ul style="list-style-type: none"> • Standalone on ASM or File Systems [2, 3] • RAC on ASM [2, 3, 4, 5] <p>Oracle 18c configured as:</p> <ul style="list-style-type: none"> • Standalone or RAC on ACFS [2, 3, 4, 5, 16, 21] <p>Oracle 19c configured as:</p> <ul style="list-style-type: none"> • Standalone only [17] 	<p>Physical [6]</p>	<ul style="list-style-type: none"> • AIX 6.1 TL9+ [8] • AIX 7.1+ [8] • AIX 7.2+ [8, 21] 	<ul style="list-style-type: none"> • Fibre Channel [18] • iSCSI [18] • NFS 	<ul style="list-style-type: none"> • HPE Nimble Storage 5.2 and later [18, 19, and 20] • IBM Spectrum Virtualize Software 7.3 and later/8.1.2 and later: <ul style="list-style-type: none"> ◦ IBM SAN Volume Controller [22] ◦ IBM Storwize [22] ◦ IBM FlashSystem V9000 and 9100 [22] • IBM Spectrum Accelerate 11.5.3 and later: <ul style="list-style-type: none"> ◦ IBM FlashSystem A9000/A9000R [22] ◦ IBM XIV storage systems [22] • NetApp ONTAP storage systems running the following versions [14][15]: <ul style="list-style-type: none"> ◦ Data ONTAP 8.1.0, 8.2.0 and later ◦ Data ONTAP 9.x ◦ Clustered Data ONTAP 8.1, 8.2, 8.3, 9.4, 9.5, 9.6 and later • Pure Storage running Pure APIs 1.5 and later: <ul style="list-style-type: none"> ◦ FlashArray//c

		<ul style="list-style-type: none"> • Red Hat Enterprise Linux / Centos / Oracle 6.5+ [9] • Red Hat Enterprise Linux / Centos / Oracle 7.0+ [9] • SUSE Linux Enterprise Server 11.0 SP4+ [9] • SUSE Linux Enterprise Server 12.0+ [9] 	<ul style="list-style-type: none"> • Fibre Channel [18] • iSCSI [18] • NFS 	
	<p>Virtual (VMware) [6, 7, 10, 11, 12]</p>	<ul style="list-style-type: none"> • Red Hat Enterprise Linux / Centos / Oracle 6.5+ [9] • Red Hat Enterprise Linux / Centos / Oracle 7.0+ [9] • SUSE Linux Enterprise Server 11.0 SP4+ [9] • SUSE Linux Enterprise Server 12.0+ [9] 	<ul style="list-style-type: none"> • Physical RDM backed by Fibre Channel or iSCSI disks attached to ESX [18] • VMDK (dependent as well as independent disks) on VMFS datastores (Fibre Channel / iSCSI) or NFS datastores [18] • iSCSI disks directly attached to guest operating system [18] • NFS share mapped directly to the guest 	<ul style="list-style-type: none"> ◦ FlashArray//m ◦ FlashArray//x ◦ FlashArray 4xx series

[1] For Oracle 12c multitenant databases, IBM Spectrum Copy Data Management supports protection and recovery of the container database, including all pluggable databases under it. Granular recovery of specific PDBs can be performed via Instant Disk Restore recovery combined with RMAN.

[2] Standalone databases protected by IBM Spectrum Copy Data Management can be recovered to the same or another standalone server as well as to a RAC installation. When recovering from standalone to RAC, if the source database uses Automatic Storage Management, then it will be successfully recovered to all nodes in the destination cluster. If the source database uses non-ASM storage, the database will be mounted only on the first node in the destination RAC. Source disks for standalone databases restoring to a virtual RAC environment must be thick provisioned.

RAC databases protected by IBM Spectrum Copy Data Management can be recovered to the same or another RAC installation as well as to a standalone server. In order to recover a RAC database to a standalone server, the destination server must have Grid Infrastructure installed and an ASM instance must be running.

[3] Oracle Data Guard primary and secondary databases support inventory and backup operations. Oracle Data Guard databases will always be restored as primary databases without any Data Guard configurations enabled.

[4] Oracle Flex ASM is not supported.

[5] RAC database recoveries are not server pool-aware. IBM Spectrum Copy Data Management can recover databases to a RAC, but not to specific server pools.

[6] IBM Spectrum Copy Data Management supports recovering databases from a source physical server to a destination virtual server by provisioning disks as physical RDMs. Similarly, IBM Spectrum Copy Data Management can recover databases from a source virtual server that uses physical RDM to a destination physical server. However, source databases on VMDK virtual disks can only be recovered to another virtual server and not to a physical server.

[7] IBM Spectrum Copy Data Management does not support VADP-based protection of virtual Oracle servers. Oracle data must reside directly on one of the supported storage systems listed above.

[8] On AIX LPAR/VIO servers, Oracle data must reside on disks attached to the server using NPIV. Virtual SCSI disks are not supported.

[9] Linux LVM volumes containing Oracle data must use LVM version 2.02.118 or above. On SLES 11 SP4, this version of LVM may not be available through the official repositories, in which case, databases running on or recovered to SLES 11 SP4 systems must use ASM or non-LVM filesystems only.

[10] Masking and DevOps recoveries are not supported on virtual servers.

[11] See [System Requirements](#) on page 23 for supported VMware vSphere versions.

[12] Oracle servers registered as virtual must have VMware Tools installed and running.

[13] Oracle 12c multithreaded configurations are not supported.

[14] Data masking is not supported for Oracle in NetApp storage environments. Masking is not supported on source database on NFS (a copy of which will be cloned and masked) or source databases on replica copies. Instead of the default mirror copy, you must select a snapshot copy as a replication source.

[15] NetApp systems running in 7-Mode are not supported.

[16] The `ORACLE_HOME` directory may reside on ACFS. The `data` and `log` directories are required to be on ACFS. IBM Spectrum Copy Data Management does not protect the Oracle installation that is contained in the `ORACLE_HOME` directory.

[17] Oracle 19c is only supported for standalone on AIX 7.2. It is not supported using RAC.

Additional Requirements:

- Oracle database data and the flash recovery area (FRA) must reside on supported storage systems. IBM Spectrum Copy Data Management can back up archived logs to a supported storage system if they are not already on one.

- During an Instant Database Restore, there may be failures if the new name specified for the restored database is similar to an existing database only differing by numerical suffix. For clustered instances of Oracle databases, the appliance always uses global database name in the UI. During the inventory and restore processes, individual instances using the numerical suffixes of the cluster must be correlated to the global database name. The issue with this comes when, as an example, "Production12" is discovered. Is this the instance 12 of the "Production" database, or instance two (2) of the "Production1" database, or a database named "Production12."

[18] HPE Nimble Storage must be version 5.2 or later to support iSCSI and Fibre Channel.

[19] Oracle PIT (point-in-time) recovery is not supported with HPE Nimble Storage.

[20] The make permanent option is not available for HPE Nimble Storage using physical disks.

[21] Oracle ACFS on AIX 7.2 is supported. Oracle ACFS on AIX is not supported for flashcopy incremental backups and iSCSI protocol.

[22] On IBM Systems Storage, condense is run during maintenance jobs.

For Oracle RAC clustered nodes running prior to vSphere 6.0, virtual machines cannot use a virtual SCSI controller whose SCSI Bus Sharing option is set to None. This is required to ensure that IBM Spectrum Copy Data Management can hot-add shared virtual disks to the cluster nodes. For vSphere 6.0 and above, this requirement does not apply. Instead, if an existing shared SCSI controller is not found on vSphere 6.0 and above, IBM Spectrum Copy Data Management automatically enables the "multi-writer" sharing option for each shared virtual disk.

Does IBM Spectrum Copy Data Management support Oracle running on a VMware VM?

Oracle support for VMware virtual machines requires Oracle data/logs to be stored on VMDK virtual disks or Physical RDMs (pRDM). Virtual RDM disks are not supported. The VMDKs must reside on a datastore created on LUNs from supported storage systems. Similarly, the Physical RDMs must be backed by LUNs from supported storage systems.

Does IBM Spectrum Copy Data Management support Oracle Database 12c Multitenant features?

IBM Spectrum Copy Data Management supports the Oracle Database 12c R1 Multitenant option for backup or clone of a container database (CDB). Recoveries of pluggable databases (PDB) are supported through RMAN.

How can granular recovery of specific PDBs be performed?

Granular recovery of specific Pluggable Databases (PDBs) can be performed via Instant Disk Restore recovery combined with RMAN. To do so:

1. Perform an Instant Disk Restore of the Container Database (CDB) by using an Oracle Restore job in IBM Spectrum Copy Data Management. The Oracle CDB backup may already have been cataloged into RMAN when the backup was created, or you can opt to do this during the Instant Disk Restore.

2. Login to RMAN. List the copies in the catalog and identify the tag from which you want to recover. Tags for IBM Spectrum Copy Data Management-created entries are generally of the form "ECX_<timestamp>".

3. Close the existing PDB by running the following command:

```
alter pluggable database <PDB_name> close;
```

4. Recover the PDB by running the following command:

```
run {
  restore pluggable database <name> from tag '<tag_name>';
  recover pluggable database <name>;
}
```

5. Open the recovered PDB by running the following command:

```
alter pluggable database <PDB_name> open;
```

Can I back up Oracle running on any storage to a supported storage system via VADP (VM Replication job)?

No, not through IBM Spectrum Copy Data Management Oracle Backup workflow. You can leverage a VMware Backup job with pre/post script to protect Oracle database in such configuration.

Oracle Requirements

Review the following requirements and pre-requisites for registering an Oracle provider in IBM Spectrum Copy Data Management.

Software

- The *bash* and *sudo* packages must be installed. *Sudo* must be version 1.7.6p2 or above. Run `sudo -V` to check the version.
- Python version 2.6.x or 2.7.x must be installed.
- **AIX only:** If Oracle data resides on IBM Spectrum Accelerate storage, the IBM Storage Host Attachment Kit (also known as IBM XIV Host Attachment Kit) must be installed on the Oracle server.
- **RHEL/OEL/CentOS 6.x only:** Ensure the *util-linux-ng* package is up-to-date by running `yum update util-linux-ng`. Depending on your version or distribution, the package may be named *util-linux*.
- **RHEL/OEL/CentOS 7.3 and above:** A required Perl module, `Digest::MD5`, is not installed by default. Install the module by running `yum install perl-Digest-MD5`.
- **Linux only:** If Oracle data resides on LVM volumes, ensure the LVM version is 2.0.2.118 or later. Run `lvm version` to check the version and run `yum update lvm2` to update the package if necessary.

- **Linux only:** If Oracle data resides on LVM volumes, the *lvm2-lvmetad* service must be disabled as it can interfere with IBM Spectrum Copy Data Management's ability to mount and resignature volume group snapshots/clones.

Run the following commands to stop and disable the service:

```
$ systemctl stop lvm2-lvmetad
$ systemctl disable lvm2-lvmetad
```

Additionally, disable *lvmetad* in the LVM config file. Edit the file */etc/lvm/lvm.conf* and set:

```
use_lvmetad = 0
```

Connectivity

- The SSH service must be running on port 22 on the server and any firewalls must be configured to allow IBM Spectrum Copy Data Management to connect to the server using SSH. The SFTP subsystem for SSH must also be enabled.
- The server can be registered using a DNS name or IP address. DNS names must be resolvable by IBM Spectrum Copy Data Management.
- When registering Oracle RAC nodes, register each node using its physical IP or name. Do not use a virtual name or Single Client Access Name (SCAN).
- In order to mount clones/copies of Oracle data, IBM Spectrum Copy Data Management automatically maps and unmaps LUNs to the Oracle servers. Each server must be preconfigured to connect to the relevant storage systems at that site.
 - For Fibre Channel, the appropriate zoning must be configured beforehand.
 - For iSCSI, the Oracle servers must be configured beforehand to discover and log in to the targets on the storage servers.

Authentication

- The Oracle server must be registered in IBM Spectrum Copy Data Management using an operating system user that exists on the Oracle server (referred to as "IBM Spectrum Copy Data Management agent user" for the rest of this topic).
- During registration you must provide either a password or a private SSH key that IBM Spectrum Copy Data Management will use to log in to the server.
- For password-based authentication ensure the password is correctly configured and that the user can log in without facing any other prompts, such as prompts to reset the password.
- For key-based authentication ensure the public SSH key is placed in the appropriate *authorized_keys* file for the IBM Spectrum Copy Data Management agent user.
 - Typically, the file is located at */home/<username>/.ssh/authorized_keys*
 - Typically, the *.ssh* directory and all files under it must have their permissions set to 600.

Privileges

The IBM Spectrum Copy Data Management agent user must have the following privileges:

- Privileges to run commands as root and other users using sudo. IBM Spectrum Copy Data Management requires this for various tasks such as discovering storage layouts and mounting and unmounting disks.
 - The *sudoers* configuration must allow the IBM Spectrum Copy Data Management agent user to run commands without a password.
 - The *!requiretty* setting must be set.
 - The *ENV_KEEP* setting must allow the *ORACLE_HOME* and *ORACLE_SID* environment variables to be retained.
- Privileges to read the Oracle inventory. IBM Spectrum Copy Data Management requires this to discover and collect information about Oracle homes and databases.
 - To achieve this, the IBM Spectrum Copy Data Management agent user must belong to the Oracle inventory group, typically named *oinstall*.
- SYSDBA privileges for database instances. IBM Spectrum Copy Data Management needs to perform database tasks like querying instance details, hot backup, RMAN cataloging, as well as starting/stopping instances during recovery.
 - To achieve this, the IBM Spectrum Copy Data Management agent user must belong to the OSDBA operating system group, typically named *dba*.
 - In the case of multiple Oracle homes each with a different OSDBA group, the IBM Spectrum Copy Data Management agent user must belong to each group.
- SYSASM privileges, if Automatic Storage Management (ASM) is installed. IBM Spectrum Copy Data Management needs to perform storage tasks like querying ASM disk information, as well as renaming, mounting, and unmounting diskgroups.
 - To achieve this, the IBM Spectrum Copy Data Management agent user must belong to the OSASM operating system group, typically named *asmadmin*.
- Shell user limits for the IBM Spectrum Copy Data Management agent user must be the same as those for the user that owns the Oracle home, typically named *oracle*. Refer to Oracle documentation for requirements and instructions on setting shell limits. Run `ulimit -a` as both the *oracle* user and the IBM Spectrum Copy Data Management agent user and ensure their settings are identical.

For examples on creating a new user with the necessary privileges, see [Sample Configuration of an IBM Spectrum Copy Data Management Agent User](#) on page 582.

Oracle Home Discovery for Symbolic Links

IBM Spectrum Copy Data Management discovers Oracle installations and databases by looking through the files

`/etc/oraInst.loc` and `/etc/oratab`, we well as the list of running Oracle processes. When connecting to a database instance to discover its properties, IBM Spectrum Copy Data Management connects by setting the `ORACLE_HOME` environment variable based on the path that was auto-discovered. In some cases, when Oracle databases are used in combination with SAP software for instance, a symbolic link to the Oracle Home path may be created. This link path can then be used by the database administrator to connect to the database.

Consider the following example where the real Oracle Home path can be

`/u01/app/oracle/product/1210/dbhome_1`, while the symbolic link that points to it is `/oracle/PRODDB/1210`.

```
$ ls -l /oracle/PRODB
lrwxrwxrwx. 1 oracle oinstall 1210 -> /u01/app/oracle/product/1210/dbhome_1
```

When IBM Spectrum Copy Data Management connects to the instance using the real, auto-discovered path, queries against the database can fail due to the `ORACLE_HOME` environment variable being incorrectly set. To override the auto-discovered path and force IBM Spectrum Copy Data Management to use the correct symbolic link path, a configuration file must be defined on each Oracle server where symbolic links are in use.

Create or edit if it already exists a file named `/etc/guestapps_oraHomes.conf`. Insert into this file these entries:

```
# Define one entry per line.
# Each entry must be in the form: <instanceName> = <oraHomePath>
# Lines beginning with the '#' character are ignored.
PRODDB = /oracle/PRODDB/1210
TESTDB = /oracle/TESTDB/1210
```

To ensure that the configuration file is readable by the agent user, run the following command:

```
chmod 644 /etc/guestapps_oraHomes.conf
```

Database Discovery

IBM Spectrum Copy Data Management discovers Oracle installations and databases by looking through the files `/etc/oraInst.loc` and `/etc/oratab`, as well as the list of running Oracle processes. If the files are not present in their default location, the "locate" utility must be installed on the system so that IBM Spectrum Copy Data Management can search for alternate locations of these files.

IBM Spectrum Copy Data Management discovers databases and their storage layouts by connecting to running instances and querying the locations of their datafiles, log files, etc. In order for IBM Spectrum Copy Data Management to correctly discover databases during cataloging and copy operations, databases must be in "MOUNTED," "READ ONLY," or "READ WRITE" mode. IBM Spectrum Copy Data Management cannot discover or protect database instances that are shut down.

Databases must be started using a server parameter file (spfile). IBM Spectrum Copy Data Management does not support copy operations for databases that are started using a text-based parameter file (pfile).

ASM Disk Discovery

When IBM Spectrum Copy Data Management mounts snapshots/clones of ASM disks, it configures the disks to set the appropriate permissions required to make them discoverable by ASM:

- The disk owner and group are set to the owner of the Grid installation and the OSASM group respectively. These are typically *grid* and *asmadmin*. IBM Spectrum Copy Data Management automatically discovers the appropriate owner and group information on each server.
- The disk permissions are set to 660.

Additionally, IBM Spectrum Copy Data Management creates aliases/symbolic links with names that follow a consistent pattern. To ensure that ASM is able to discover the disks mapped by IBM Spectrum Copy Data Management, you must update the *ASM_DISKSTRING* parameter to add this pattern.

Linux:

IBM Spectrum Copy Data Management creates *udev* rules for each disk to set the appropriate ownership and permissions. The *udev* rules also create symbolic links of the form */dev/ecx-asmdisk/<diskId>* that point to the appropriate device under */dev*.

To ensure the disks are discoverable by ASM, add the following pattern to your existing *ASM_DISKSTRING*:
*/dev/ecx-asmdisk/**

AIX:

IBM Spectrum Copy Data Management creates a device node (using *mknod*) of the form */dev/ecx_asm<diskId>* that points to the appropriate *hdisk* under */dev*. IBM Spectrum Copy Data Management also sets the appropriate ownership and permissions for this new device.

To ensure that the disks are discoverable by ASM, add the following pattern to your existing *ASM_DISKSTRING*: */dev/ecx_asm**

Notes:

- If the existing value of the *ASM_DISKSTRING* is empty, you may have to first set it to an appropriate value that matches all existing disks, then append the value above.
- If the existing value of the *ASM_DISKSTRING* is broad enough to discover all disks (for example, */dev/**), you may not need to update it.
- Refer to Oracle documentation for details about retrieving and modifying the *ASM_DISKSTRING* parameter.

Sample Configuration of an IBM Spectrum Copy Data Management Agent User

The commands below are examples for creating and configuring an operating system user that IBM Spectrum Copy Data Management will use to log in to the Oracle server. The command syntax may vary depending on your operating system type and version.

- Create the user that will be designated as the IBM Spectrum Copy Data Management agent user: `useradd -m cdmagent`
- Set a password if using password-based authentication: `passwd cdmagent`

- If using key-based authentication, place the public key in `/home/cdmagent/.ssh/authorized_keys`, or the appropriate file depending on your `sshd` configuration, and ensure the correct ownership and permissions are set, such as:

```
chown -R cdmagent:cdmagent /home/cdmagent/.ssh
```

```
chmod 700 /home/cdmagent/.ssh
```

```
chmod 600 /home/cdmagent/.ssh/authorized_keys
```

- Add the user to the Oracle installation and OSDBA group:

```
usermod -a -G oinstall,dba cdmagent
```
- If ASM is in use, also add the user to the OSASM group:

```
usermod -a -G asmadmin cdmagent
```

Note: If on AIX, the append argument (`-a`) should be omitted when using the `usermod` command.

- Place the following lines at the end of your `sudoers` configuration file, typically `/etc/sudoers`. If your existing `sudoers` file is configured to import configuration from another directory (for example, `/etc/sudoers.d`), you can also place the lines in a new file in that directory:

```
Defaults:cdmagent !requiretty
```

```
Defaults:cdmagent env_keep+="ORACLE_HOME"
```

```
Defaults:cdmagent env_keep+="ORACLE_SID"
```

```
cdmagent ALL=(ALL) NOPASSWD:ALL
```

Microsoft SQL Server Support FAQ

What is SQL Server CDM? How does it help solve my challenges?

Database administrators working extensively with SQL Server are challenged when faced with mission-critical use cases such as Backup, Recovery, DevOps, and Business Analytics. This is especially true given that their SQL Server databases have expanded in size and number over time, and that the databases need to be up and running 24x7x365.

SQL Server DBAs struggle with the following:

- Backups are slow, complex and need constant management
- Backup process slows down the production servers
- Recoveries are slow and complex
- Repurposing App consistent backups (clones) for DevOps and Business Analytics is slow, complex and storage inefficient
- Lack of automation exists for producing quick and secure clones required to accelerate DevOps
- Copy sprawl problems occur due to no central catalog of copies
- Unable to meet organization's stringent RPO and RTO requirements

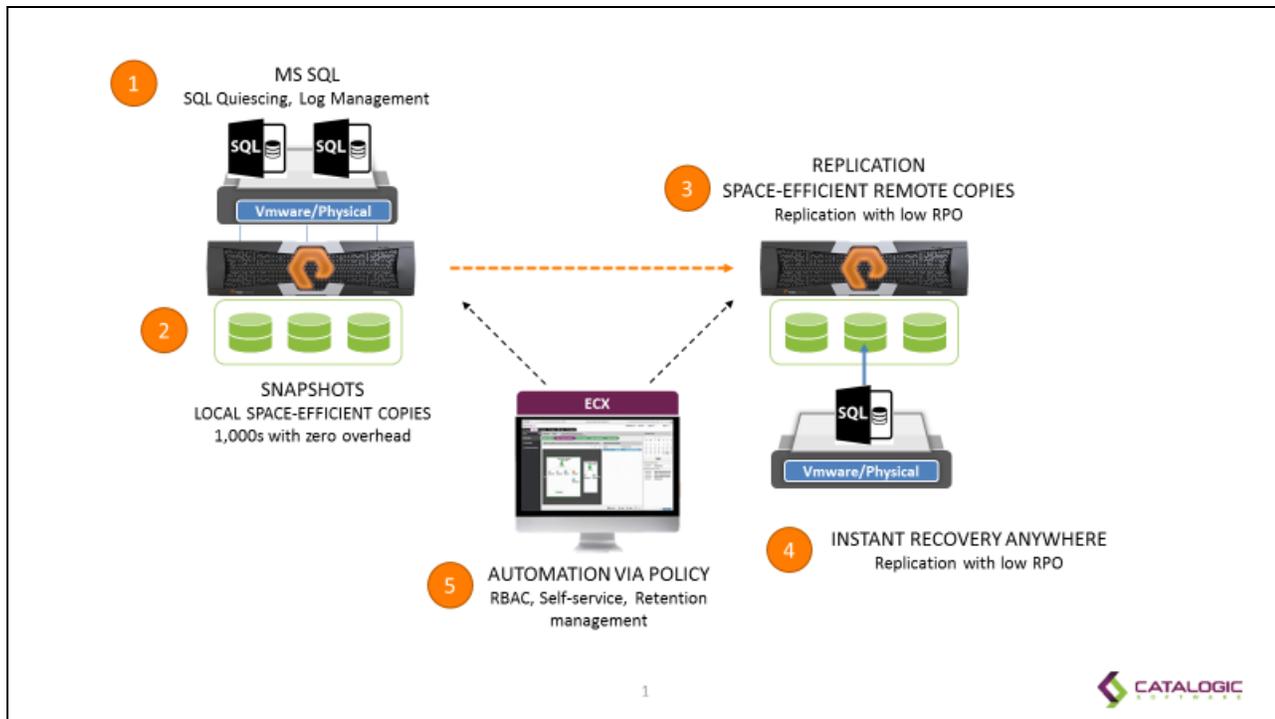
IBM Spectrum Copy Data Management simplifies SQL Server copy management by enabling administrators to orchestrate application-consistent copy creation, cloning and recovery in minutes, instead of hours or days. IBM Spectrum Copy Data Management copy management leverages the advanced snapshot and replication features of the underlying storage platform to rapidly create, replicate, clone, and restore copies of SQL Server databases in the most efficient way possible, in both time and space. IBM Spectrum Copy Data Management enables you to focus on the backup and restore requirements of your business rather than the technical details of the underlying storage platforms.

IBM Spectrum Copy Data Management is an intelligent copy data management solution that delivers end-to-end automation, orchestration, and self-service functionality for your SQL Server environment through a comprehensive and scalable catalog. With the self-service features of IBM Spectrum Copy Data Management, your users are empowered to create clones on demand, freeing DBAs, while at the same time offering the advanced recovery features needed for SQL Server environments.

IBM Spectrum Copy Data Management SQL Server Copy Data Management solution supports the following SQL Server deployment modes running on VMware virtual machines or physical servers:

- Standalone SQL Server – Databases running on a single server
- SQL Server Failover Cluster – SQL Server instances running on Windows Server Failover Clusters using Shared storage
- SQL Server Always On – Primary and secondary databases in Availability group configured across clusters of servers

IBM Spectrum Copy Data Management SQL Server CDM Solution Architecture



Deployment and Registration

Do I need to deploy any additional agents to protect SQL Server standalone, Failover Cluster or AlwaysON configuration?

IBM Spectrum Copy Data Management for SQL Server is delivered as a VMware OVA that is easily deployed on demand in a matter of minutes. Once deployed, you simply register your SQL Servers with appropriate credentials and then let IBM Spectrum Copy Data Management discover the rest. IBM Spectrum Copy Data Management eliminates the complexity of manually deploying and maintaining application agents on SQL Servers. A lightweight application-aware agent is automatically injected and updated to the required SQL Servers on demand.

SQL Server Backup workflow

Application-consistent SQL Server backup (local and remote) - Step by step

IBM Spectrum Copy Data Management auto-discovers databases and enables copies only of eligible databases. To be eligible for IBM Spectrum Copy Data Management backup, the SQL Server database needs to be residing on a supported storage platform. With IBM Spectrum Copy Data Management, application owners do not need to be concerned about storage infrastructure.

A typical SQL Server database backup creation workflow consists of following steps:

- Auto-inject lightweight agent into SQL Server node running database instance
- Discover storage volume mapping to selected SQL Server database(s) and logs

- Place SQL Server database in hot backup mode via VMware Snapshot/VSS Snapshot (App consistent)
- Automatically create consistency group for related storage volumes (only on physical servers)
- Create application-consistent backup (VADP backup and/or Storage snapshot)
- Take SQL Server database out of hot backup mode (Delete VM snapshot/VSS snapshot)
- Optionally create log copies with lowest RPO possible
- Catalog SQL Server database backups in catalog
- Optionally replicate application-consistent copy to remote location leveraging storage replication feature

IBM Spectrum Copy Data Management creates and uses in-place copies, so no data is physically moved. IBM Spectrum Copy Data Management generated application-consistent copies are both space and time efficient. With the same ease, a DBA can automate the creation of remote copies for disaster recovery use cases.

Does SQL Server solution leverage the storage consistency group feature?

The storage consistency group feature allows storage administrators to take a snapshot of database applications where the data is spread across multiple volumes to maintain consistency across all volumes.

In a typical SQL Server Database, the data is spread across different volumes for better IO performance and availability. On Physical servers, IBM Spectrum Copy Data Management SQL Server application-consistent copy creation ensures that appropriate consistency groups are automatically created to maintain consistency across all related volumes. IBM Spectrum Copy Data Management SQL Server backup on VM relies on VMware snapshots and doesn't need to leverage storage consistency group feature.

What level of Application selection granularity is supported for SQL Server Backup jobs?

IBM Spectrum Copy Data Management SQL Server backup job definition supports copy selection at the following levels:

- One or more SQL Server Instances for Standalone SQL Server/Failover Cluster
- One or more Availability groups for AlwaysON
- One or more Databases for Standalone/Failover Cluster and SQL Server AlwaysON

Can I restore a database to an original instance and overwrite existing database in a single step?

Yes, use the Overwrite existing database option in the Application restore job definition.

Will IBM Spectrum Copy Data Management auto discover newly added SQL Server instances in a Standalone SQL Server and automatically protect it?

No. IBM Spectrum Copy Data Management will auto discover and present newly added SQL Server instances in the Backup job but you must explicitly select newly added SQL server instances for protection. Discovery of new SQL Server instances occurs as part of a regularly scheduled Application inventory job.

Will IBM Spectrum Copy Data Management auto discover newly added databases and automatically protect it?

Yes, if you select at Availability group level protection, IBM Spectrum Copy Data Management will auto discover newly added databases in selected availability group and protect it automatically during next job instance run. Discovery of new SQL Server instances and database occurs as part of regularly scheduled Application Inventory job.

Does IBM Spectrum Copy Data Management backup primary databases or secondary databases in SQL AlwaysOn?

IBM Spectrum Copy Data Management backs up only primary databases across the SQL AlwaysOn cluster.

Do SQL Server databases and logs need to be on supported storage for IBM Spectrum Copy Data Management CDM?

IBM Spectrum Copy Data Management also supports protection of SQL Server running on VMware VM configured on any storage that can be protected to supported storage systems via VM Replication. SQL Server running on physical servers require the database and logs to be on supported storage.

Does IBM Spectrum Copy Data Management perform full backups of databases?

IBM Spectrum Copy Data Management backups of SQL Server databases are always VSS COPY type backups.

SQL Server log management

Does IBM Spectrum Copy Data Management support Transaction log backup and log management?

Every SQL Server database has a transaction log that records all transactions and the database modifications made by each transaction. The transaction log must be truncated on a regular basis to keep it from filling up.

IBM Spectrum Copy Data Management provides you with an option to back up transaction log files. IBM Spectrum Copy Data Management supports log backup at a specified frequency. You can select one or more databases for log backup in a single backup job definition. Log destination can be specified as a single universal mount point or separate destination mount point for each database. Specified log backup destination path must already exist and must reside on supported storage system. If multiple databases are selected for backup, then each of the servers hosting the database must have their Destination directory set individually.

Does IBM Spectrum Copy Data Management support truncation of database logs?

Yes, IBM Spectrum Copy Data Management will automatically truncate log post log backups of databases that it backs up. If database logs are not backed up with IBM Spectrum Copy Data Management, its logs are not truncated by IBM Spectrum Copy Data Management and must be managed separately.

Can I specify a retention period for backed up transaction logs?

No. Log backup retention and auto-deletion is planned for a future release.

I am backing up transaction logs with IBM Spectrum Copy Data Management, but I don't want it to truncate logs. Can I control this behavior?

No. This will be enhanced in future release of IBM Spectrum Copy Data Management.

Pre and Post Scripts

Does IBM Spectrum Copy Data Management support pre/post scripts for Application Database Backup jobs?

Yes, IBM Spectrum Copy Data Management supports job-level pre/post scripts and job-level pre/post Snapshot scripts to enable further customization.

Job-level prescripts and postscripts are scripts that can be run before or after a job runs.

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot subpolicy runs. (Please refer to pre/post script topic in the IBM Spectrum Copy Data Management User's Guide for details.)

Data masking

Does IBM Spectrum Copy Data Management SQL Server solution offer Data Masking integration with third party masking tools?

A concern for security officers in any organization is that of keeping confidential information locked down, even internally. Data masking is used to hide confidential data, by replacing it with fictitious data, when making data copies for DevTest or other use cases. It prevents leakage of sensitive data in non-production databases via static data masking [SDM], and production data in transit via dynamic data masking [DDM].

The following Data Masking integration features will be available in a future release.

IBM Spectrum Copy Data Management will include integrated data masking workflows with the ability to leverage third party masking tools. Traditionally, data masking is difficult, slow, and storage-consuming, but with IBM Spectrum Copy Data Management it will be easily integrated into the SQL Server backup workflow, allowing creation of masked copies at a specified frequency. Masked copies are automatically marked in the catalog. Access to secure copies is managed by the administrator by leveraging the application-level RBAC.

In addition, SQL Server will enable you to leverage the Dynamic Data Masking feature of SQL Server 2016.

Is sample masking script provided with IBM Spectrum Copy Data Management SQL Server solution?

A sample data masking script can be provided upon request. A sample masking script demonstrates data masking integration with built-in Dynamic Data masking script of SQL Server 2016. This feature will be available in a future release.

SQL Server Restore Workflow

Can I leverage SQL Server database clones for multiple use?

Limitations of current tools and approaches:

- Database cloning requires action by DBAs and is gated by process

- QA relies on DBAs for cloning the databases for functional testing
- The database cloning is gated by processes (space requisition, approvals, etc.)
- Database cloning is not time- or storage-efficient
- Usage of common cloning tools or custom scripts creates full copy requiring large amounts of additional storage
- Creating full copies is slow

IBM Spectrum Copy Data Management solves these challenges with simple, automated end-to-end clone lifecycle management:

- Self-service access to secure clones by QA team eliminates administrative and process bottlenecks
- IBM Spectrum Copy Data Management enables rapid database clones that are both time- and space-efficient
- Provision clones in minutes regardless of its size
- Leverages underlying storage snapshots for space efficiency
- IBM Spectrum Copy Data Management promotes standardization and governance through centralized catalog, granular RBAC, and automated policies

Your SQL Server clones can be utilized and consumed instantly – for whatever your use case -- through IBM Spectrum Copy Data Management “Instant Disk Restore” jobs. IBM Spectrum Copy Data Management catalogs and tracks all cloned instances. Instant Disk Restore can leverage iSCSI or FC protocol to provide immediate mount of LUNs without transferring data.

Can I create an Instant Clone of a SQL Server database for DevOps and Business Analytics?

IBM Spectrum Copy Data Management provides automated workflows to create instant clones of SQL Server database regardless of its size.

- Instantly create database clones from any of the copies in the IBM Spectrum Copy Data Management inventory, at local or remote locations, to accelerate Business Analytics.
- Enable and accelerate DevOps by providing Instant Disk Restore to secure clones of databases to appropriate users via application-level RBAC.

Then, when your TestDev, DevOps, or research/analytics work is completed, you can save the clone to more permanent storage or simply tear it down.

What is the granularity of database recovery supported by IBM Spectrum Copy Data Management?

Supported recoveries for standalone or AlwaysOn:

- Database can be recovered to point of snapshots to original or new instance (Instant Disk Restore)
- Database can be recovered to point in time leveraging backed up transaction logs (Instant recovery) to original or new instance
- Database can be recovered using new name to original or new instance
- You can select one or more databases in a single restore job definition.
- Each selected database in a Restore job definition can have separate destination specification
- Databases are always recovered in online mode
- Database can be recovered from standalone instance to AlwaysON Availability group
- Database from AlwaysON Availability can be recovered to standalone instance
- Database running on older version can be recovered to instance running same or newer version.

What granularity of Point in Time recovery is supported?

IBM Spectrum Copy Data Management enables database recovery to a specific point in time, allowing you to:

- Restore to the state just before the point of failure
- Restore multiple databases to a consistent time

Can I restore a Database to a Transaction Mark?

No. This will be enhanced in a future release of the product.

Does IBM Spectrum Copy Data Management support recovering a database in online mode?

Yes, Instant Disk Restore or Instant Database Restore recovers databases in online mode.

Does IBM Spectrum Copy Data Management support recovering databases in an offline state (norecovery)?

No. This will be enhanced in a future release of the product.

Does IBM Spectrum Copy Data Management support recovering databases in a standby/read-only state (standby)?

Yes, IBM Spectrum Copy Data Management provides an application option to control this behavior.

Roll back uncommitted transactions and leave the database ready to use

- Select this option to restore the database to an online state. If selected, additional transaction logs cannot be restored. If deselected, uncommitted transactions are not rolled back, leaving the database non-operational. Additional transaction logs can then be restored.

Does IBM Spectrum Copy Data Management support recovering database with Restricted Access?

No. This will be enhanced in a future release of the product.

Does IBM Spectrum Copy Data Management support restoring only logs so that they can be applied to a standby database?

No, not from the IBM Spectrum Copy Data Management application restore workflow. A user can easily access the transaction log backup location from the SQL server and perform this outside of the product.

Does IBM Spectrum Copy Data Management support out-of-place restore?

Out-of-place restore is used to relocate a database file to a new location:

- Copying/moving a database to a different location on a same SQL Server instance
- Copying a Database to a different SQL Server Instance at a different location

This feature will be enhanced in a future release of the product.

Where are SQL specific and IBM Spectrum Copy Data Management specific logs if errors occur?

All required logs (IBM Spectrum Copy Data Management and application) are collected as part of the current log collection functionality. There should be no need to manually obtain SQL application logs from within SQL Server VMs.

How do I refresh? How to promote to Production?

All database recovery operations can leverage Instant mode (Test) and then can either be deleted or promoted to permanent mode via workflow control. This behavior can be controlled via the Make Permanent job option.

- Enabled - Always make permanent
- Disabled - Never make permanent
- User election - Allows the user to select Make Permanent or Cleanup when the job session is pending

Does IBM Spectrum Copy Data Management use existing hardware providers for physical SQL backups?

No. IBM Spectrum Copy Data Management automatically deploys its own VSS HW provider service for SQL Server running on physical servers. It is automatically started on demand during SQL Server Backup jobs. At the completion of the backup job, the VSS HW provider service is automatically stopped.

When IBM Spectrum Copy Data Management protects a SQL VM with pRDM, can it restore a database back to the original node as a pRDM?

Currently, a SQL VM with pRDM must be registered as Physical in IBM Spectrum Copy Data Management. Hence, the restoration of that data obeys the Physical restore restrictions, which means it can only restore back to the original host via iSCSI. If the target host being restored to was registered as Virtual, then the database would be restored as a pRDM.

This functionality will be improved in future release.

Why must I choose a proxy node when performing a restore to a SQL Failover cluster?

Windows requires signatures to be unique, so when you attach a disk that has a signature equal to one that is already attached, Windows keeps the disk in “offline” mode and doesn’t read its partition table or mount its volumes. To prevent disk signature collision, during Instant Database Restore, IBM Spectrum Copy Data Management leverages Windows proxy servers to temporarily mount disks from snapshots, generate a new signature, then mount to original server.

Any Windows node with iSCSI or Fibre Channel access to the storage can be selected as a proxy server, provided that the node is not part of the original cluster. It is recommended to select a standalone virtual or physical Windows node as a proxy server.

Self Service

Does IBM Spectrum Copy Data Management support RBAC? What is the level of granularity supported for SQL Servers?

Role-based access control allows you to set the resources and permissions available to IBM Spectrum Copy Data Management accounts. Through role-based access control you can tailor IBM Spectrum Copy Data Management for individual users, giving them access to the features and providers they need.

Using IBM Spectrum Copy Data Management RBAC functionality, user can delegate IBM Spectrum Copy Data Management role to enable and accelerate DevOps by providing Instant Access to secure clones of databases to appropriate users via application-level RBAC. Then, when your TestDev, DevOps, or research/analytics work is completed, you can save the clone to more permanent storage or simply tear it down.

Can developers access IBM Spectrum Copy Data Management operations using command line or APIs?

A rich set of REST APIs are provided to enable full access to IBM Spectrum Copy Data Management functionalities for further customization. Please refer to Product Marketplace for full API documentation and a facility that allows to you to experiment with IBM Spectrum Copy Data Management APIs.

System Requirements

What SQL Server versions are supported and on what Windows OS? What are supported Storage Systems for Microsoft SQL Server?

The Microsoft Server SQL Support Matrix is below.

Microsoft SQL Server Configuration Requirements

Database Versions/Types and Operating Systems	Server Types	Storage Systems	Storage Configuration
<ul style="list-style-type: none"> SQL 2012 on Windows Server 2012 R2 [1] SQL 2012 SP2 on Windows 	Physical [7] [8] [9] [10] [11]	<ul style="list-style-type: none"> HPE Nimble Storage 5.2 and later [12, 13, and 14] IBM Spectrum Accelerate 11.5.3 and later: <ul style="list-style-type: none"> IBM FlashSystem A9000/A9000R [15] IBM XIV storage systems [15] 	<ul style="list-style-type: none"> Fibre Channel [12] iSCSI [12]

<p>Server 2016</p> <ul style="list-style-type: none"> • SQL 2014 on Windows Server 2012 R2 [1] • SQL 2014 on Windows Server 2016 		<ul style="list-style-type: none"> • IBM Spectrum Virtualize Software 7.3 and later/8.1.2 and later: <ul style="list-style-type: none"> ◦ IBM SAN Volume Controller [15] ◦ IBM Storwize [15] ◦ IBM FlashSystem V9000 and 9100 [15] • Pure Storage running Pure APIs 1.5 and above: <ul style="list-style-type: none"> ◦ FlashArray//c ◦ FlashArray//m ◦ FlashArray//x ◦ FlashArray 4xx series 	
<ul style="list-style-type: none"> • SQL 2016 on Windows Server 2012 R2 [1] • SQL 2016 on Windows Server 2016 • SQL 2017 on Windows Server 2019 <p>Standalone, SQL Server Failover Clustering, and AlwaysOn configurations are supported for SQL 2012, SQL 2014, and SQL 2016.[2]</p> <p>Note: It is highly recommended to install the latest SQL Server patches and updates in your environment.</p>	<p>Virtual (VMware) [3]</p>	<ul style="list-style-type: none"> • DellEMC Unity <ul style="list-style-type: none"> ◦ EMC Unity 300, 400, 500, 600 (All-Flash and Hybrid Flash) ◦ EMC UnityVSA ◦ EMC VNXe 1600 running version 3.1.3 + ◦ EMC VNXe 3200 running version 3.1.1 + • IBM Spectrum Accelerate 11.5.3 and later: <ul style="list-style-type: none"> ◦ IBM FlashSystem A9000/A9000R [15] ◦ IBM XIV storage systems [15] • IBM Spectrum Virtualize Software 7.3 and later/8.1.2 and later: <ul style="list-style-type: none"> ◦ IBM SAN Volume Controller [15] ◦ IBM Storwize [15] ◦ IBM FlashSystem V9000 and 9100 [15] • NetApp ONTAP storage systems running the following versions: <ul style="list-style-type: none"> ◦ Data ONTAP 8.1.0, 8.2.0 and later in 7-mode ◦ Data ONTAP 9.x ◦ Clustered Data ONTAP 8.1, 8.2, 8.3, 9.4 and later 	<ul style="list-style-type: none"> • iSCSI disks attached directly to guest operating system [4, 12] • pRDM [4] • vRDM [5] • VMDK [6]

		<ul style="list-style-type: none"> • Pure Storage running Pure APIs 1.5 and above: <ul style="list-style-type: none"> ◦ FlashArray//c ◦ FlashArray//m ◦ FlashArray//x ◦ FlashArray 4xx series <p>Note: VM Replication restore jobs can be run to store off-host copies on the storage systems listed above.</p>	
--	--	--	--

[1] Windows Remote Shell (WinRM) must be enabled.

[2] Note that Clustered Shared Volumes (CSV) are not supported.

[3] See System Requirements for supported VMware vSphere versions.

[4] Select the Physical provider type when registering the provider in IBM Spectrum Copy Data Management. Recoveries require direct access to storage. Note that NetApp ONTAP and DellEMC storage systems are not supported.

[5] vRDMs are supported through VM Replication jobs.

[6] Independent disks are supported only if the underlying storage utilizes supported storage systems. Register the SQL resource as Physical when configuring the provider in IBM Spectrum Copy Data Management. Note that independent disks do not allow snapshots to be taken in VMware virtual scenarios. The above listed IBM Spectrum Accelerate, IBM Spectrum Virtualize, and Pure Storage FlashArrays are supported for physical registration.

[7] When registering physical SQL servers it is recommended to register via the DNS server. The IBM Spectrum Copy Data Management appliance must be resolvable and route-able by the DNS server; the physical SQL server will communicate back to IBM Spectrum Copy Data Management through DNS.

[8] Recovery for target servers registered as Physical provider types requires direct access to storage.

[9] Any Windows node with iSCSI or Fibre Channel access to the storage can be selected as a proxy server, provided that the node is not part of the original cluster. It is recommended to select a standalone virtual or physical Windows node as a proxy server.

[10] For physical SQL servers you must allow outgoing connections to port 8443 on the IBM Spectrum Copy Data Management appliance from the SQL server.

[11] Dynamic disks are not supported.

[12] HPE Nimble Storage must be version 5.2 or later to support iSCSI and Fibre Channel.

[13] The make permanent option is not available for HPE Nimble Storage using physical disks.

[14] SQL PIT (point-in-time) recovery is not supported with HPE Nimble Storage.

SQL servers residing on any storage can also be protected to supported storage systems through VM Replication jobs.

For both physical and virtual SQL environments, point-in-time recoveries beyond the last snapshot taken are incompatible with workflows utilizing more than one Site. In a virtual environment, the SQL server, associated vCenter, and storage must be registered to the same site. In a physical environment, the SQL server and storage must be registered to the same site.

[15] On IBM Systems Storage, condense is run during maintenance jobs.

What are Environment and permission requirements for SQL Server solution?

Note the following Microsoft environmental requirements:

- Windows Remote Shell (WinRM) must be enabled
- The SQL user must enable the public and sysadmin SQL permissions.
- The user identity must have sufficient rights to install and start the IBM Spectrum Copy Data Management Tools Service on the virtual machine node. This includes "Log on as a service" rights. For more information about the "Log on as a service" right, see <https://technet.microsoft.com/en-us/library/cc794944.aspx>.
- The fully qualified domain name must be resolvable and route-able from the IBM Spectrum Copy Data Management appliance
- The virtual machine node DNS name must be resolvable and route-able from the IBM Spectrum Copy Data Management appliance
- The VMGuest version must be current
- VMware Tools must be installed on the virtual machine node

Does IBM Spectrum Copy Data Management support SQL Server 2016 running on Windows 2016?

Yes. See the matrix above.

Does IBM Spectrum Copy Data Management support SQL Server configured as Physical RDMs, or Independent disks?

Yes. See footnotes 4 and 6 in the matrix above.

Does IBM Spectrum Copy Data Management support SQL Server configured as Virtual RDMs?

Yes. For limitations see footnote 5 in the matrix above.

Does IBM Spectrum Copy Data Management support SQL Server running on physical machine(s)?

Yes. See the matrix above.

Are there additional requirements for SQL support in IBM Spectrum Copy Data Management?

SQL Support for VMware Virtual Machines

UUID must be enabled to perform Microsoft SQL-based backup functions. To enable, power off the guest machine through the vSphere client, then select the guest and click **Edit Settings**. Select **Options**, then

General under the Advanced section. Select **Configuration Parameters...**, then find the `disk.EnableUUID` parameter. If set to FALSE, change the value to TRUE. If the parameter is not available, add it by clicking **Add Row**, set the value to TRUE, then power on the guest.

The virtual machine must use SCSI disks only, dynamic disks are not supported.

The latest VMware Tools must be installed on the virtual machine node.

In-Memory OLTP Requirements and Limitations

In-Memory OLTP is a memory-optimized database engine used to improve database application performance, supported in SQL 2014 and 2016. Note the following IBM Spectrum Copy Data Management requirements and limitations for In-Memory OLTP usage:

- The maximum restore file path must be less than 256 characters, which is a SQL requirement. If the original path exceeds this length, consider using a customized restore file path to reduce the length.
- The metadata that can be restored is subject to VSS and SQL restore capabilities.

SQL Server Failover Clustering Requirements for Windows Server 2008 R2

The Failover Cluster Manager Snap-In must be imported and configured before running IBM Spectrum Copy Data Management Backup and Restore jobs. To import, run Windows PowerShell in Windows Server 2008 R2 and enter the following command: `import-module failoverclusters`

For more information, see <https://technet.microsoft.com/en-us/library/ee461009.aspx>.

Registration and Authentication

Register each SQL server as a provider in IBM Spectrum Copy Data Management by name or IP address. When registering a SQL Cluster (AlwaysOn), register each node by name or IP address. The fully qualified domain name and virtual machine node DNS name must be resolvable and route-able from the IBM Spectrum Copy Data Management appliance.

The user identity must have sufficient rights to install and start the IBM Spectrum Copy Data Management Tools Service on the node. This includes "Log on as a service" rights. For more information about the "Log on as a service" right, see <https://technet.microsoft.com/en-us/library/cc794944.aspx>.

The default security policy uses the Windows NTLM protocol, and the user identity format follows the default domain\Name format.

Kerberos Requirements

Kerberos-based authentication can be enabled through a configuration file on the IBM Spectrum Copy Data Management appliance. This will override the default Windows NTLM protocol.

For Kerberos-based authentication only, the user identity must be specified in the `username@FQDN` format. The username must be able to authenticate using the registered password to obtain a ticket-granting ticket (TGT) from the key distribution center (KDC) on the domain specified by the fully qualified domain name.

Kerberos authentication also requires that the clock skew between the Domain Controller and the IBM Spectrum Copy Data Management appliance is less than 5 minutes. Note that the default Windows NTLM protocol is not time dependent.

Privileges

On the SQL server, the system login credential must have **public** and **sysadmin** permissions enabled, plus permission to access cluster resources in a SQL AlwaysOn environment. If one user account is used for all SQL functions, a Windows login must be enabled for the SQL server, with **public** and **sysadmin** permissions enabled.

Every SQL instance can use a specific user account to access the resources of that particular SQL instance.

Acronyms

A

AD

Active Directory

API

Application Programming Interface

AWS

Amazon Web Services

B

B

Bytes

C

CBT

Changed Block Tracking

CIFS

Common Internet File System

CSV

Comma-Separated Values

D

DHCP

Dynamic Host Configuration Protocol

DN

Distinguished Name

F

FCM

FlashCopy Manager

G

GB

Gigabytes

H

HTTP

Hypertext Transfer Protocol

I

IP

Internet Protocol

iSCSI

Internet Small Computer System Interface

K

KB

Kilobytes

L

LDAP

Lightweight Directory Access Protocol

LUN

Logical Unit Number

M

MB

Megabytes

N

NFS

Network File System

O

OSSV

Open Systems SnapVault

OVF

Open Virtualization Format

P

PDF

Portable Document Format

R

RBAC

Role-based access control

RDM

Raw Device Mapping

RDN

Relative Distinguished Name

RRP

Rapid Return to Production

S

SMTP

Simple Mail Transfer Protocol

SNMP

Simple Network Management Protocol

SSL

Secure Sockets Layer

SVC

SAN Volume Controller

SVM

Storage Virtual Machine

T

TB

Terabytes

U

UUID

Universally Unique Identifier

V

VADP

VMware vStorage API for Data Protection

VASA

vSphere API for Storage Awareness

VM

Virtual Machine

VMDK

Virtual Machine Disk

VMFS

Virtual Machine File System

VVOL

Virtual Volume

Terminology

The following terminology changes are implemented across the IBM Spectrum Copy Data Management user interface as of version 2.2.6.

Pre 2.2.6	2.2.6 and later
Catalog	Inventory
Copy	Backup
Hold (Job action)	Hold Schedule
Instant Access	Instant Disk Restore
Instant Virtualization	Instant VM Restore, Instant DB Restore
Pending (Job status)	Resource Active
Policy	Job, job definition
Release (Job action)	Release Schedule
Storage Workflow	SLA Policy
Use	Restore
VM Copy	VM Replication

A

account

The definition associating a resource pool with a role for a user. A user account has access to the resources and features defined in the resource pool as well as the permissions to interact with those resources and features as defined in the role.

appliance

The virtual machine containing the IBM Spectrum Copy Data Management application and Inventory, accessible through VMware vSphere client. The appliance is also referred to as the virtual appliance or virtual machine.

B

Backup job

A job that leverages Copy Data Management technology for replicating and intelligently reusing snapshots, vaults, and mirrors.

C

clone mode

A Restore job mode where virtual machines are created in a fenced network for use cases requiring permanent or long-running copies.

Copy Data Management

The ability to understand where data copies or backups are located in your IT environment, and leverage the most appropriate backup for any given use case.

D

dashboard

A feature of the Management Interface that displays an overview of your IBM Spectrum Copy Data Management environment when you first invoke IBM Spectrum Copy Data Management.

DellEMC provider

A DellEMC storage unit such as the Unity series.

F

File Analytics Reports

A category of reports that help you review your storage needs and examine your storage capacity. Use the File Analytics Reports to drill down to information about the size of files, the age of files, and type of files on your storage systems.

FlashCopy

An IBM technology that creates point-in-time snapshots.

G

Global Mirror

An IBM technology that provides data replication between two locations.

I

IBM provider

An IBM storage controller such as an SVC.

IBM Spectrum Copy Data Management

An application that delivers visibility, insight, and control into your IT environment.

IBM Spectrum Copy Data Management System Management Reports

A category of reports that offer an in-depth view of the status of your IBM Spectrum Copy Data Management configuration, including cataloged storage system information and jobs.

Instant Access

A method of copy data management that provides instant writable access to copied data. A snapshot is mapped to a target server where it can be accessed, copied, or put immediately into production use as needed.

Instant Virtualization

A method of copy data management that creates a virtual machine by directly mapping snapshots for fast startup.

Inventory

The IBM Spectrum Copy Data Management database where object metadata is stored for rapid retrieval and reporting. Cataloging objects located on a resource enables browsing, searching, analyzing, copying, and using data in IBM Spectrum Copy Data Management. The IBM Spectrum Copy Data Management Inventory resides on the appliance.

Inventory Browser

A feature enabling you to explore cataloged resources and their underlying objects and object metadata. Drill into the Inventory Browser to logically view the details of the objects underlying a storage system, virtual host, or application. The Browse and Search features provide visibility into your IT environment.

Inventory job

A job for gathering and recording objects and object metadata about specified resources.

J

Job definition

A defined set of tasks and rules that govern the execution of a job. A job definition can be applied to one or more jobs. IBM Spectrum Copy Data Management includes support for the following job types: Inventory jobs, Backup jobs, Restore jobs, Report jobs, and Script jobs.

L

LDAP provider

A server that accesses centralized data including user and authentication information.

M

management interface

The browser-based portal through which the IBM Spectrum Copy Data Management application is viewed and managed.

mirror

An efficient block-level data replication technique that produces exact replicas for disaster recovery.

N

NetApp provider

A NetApp storage system such as a FAS.

O

object

A resource that has metadata about it stored in the Inventory. Examples of objects are files, directories, qtrees, and volumes.

P

production mode

A Restore job mode which restores a resource to the production environment from secondary storage or a remote disaster recovery site.

Protection Compliance Reports

A category of reports that help ensure your data is protected through user-defined recovery point objective parameters.

provider

An object repository, physical or virtual, from which object metadata is retrieved.

Provider Browser

A feature enabling you to view a list of registered resources and their underlying resources. The Provider Browser scans the actual resource and returns native properties.

R

RBAC

A feature that allows an administrator to set the resources and permissions available to IBM Spectrum Copy Data Management user accounts.

register

The process that allows IBM Spectrum Copy Data Management to recognize a resource. The resource can subsequently be cataloged through an Inventory job.

report job

A job that runs predefined or customized reports summarizing information about objects and object metadata managed by cataloged resources. A report job is a type of

System job and provides insight into your IT environment.

resource pool

A defined set of resources, such as providers, features, and screens, that are available to a user. The resource pool is associated with a role when defining a user account.

Restore job

A job that leverages Copy Data Management technology for reusing and recovering resources from snapshots, vaults, and mirrors.

RMAN

Oracle Recovery Manager (RMAN), a command-line and Enterprise Manager-based tool, is the method preferred by Oracle DBAs for backup and recovery of Oracle databases, including maintaining an RMAN repository.

role

A defined set of actions that can be performed by a user. The role is associated with a resource pool when defining a user account.

role-based access control

A feature that allows an administrator to set the resources and permissions available to IBM Spectrum Copy Data Management user accounts.

S

schedule

A defined set of triggers and parameters that govern when a job is run.

Script job

A job that runs a set of commands on the IBM Spectrum Copy Data Management appliance. The script job is one way to extend

the functionality of IBM Spectrum Copy Data Management.

Search

A feature used to find objects on resources that are cataloged in IBM Spectrum Copy Data Management. IBM Spectrum Copy Data Management searches the entire Inventory and returns all objects with a name that matches or contains the search entry. Additionally, you can search with filters such as object name, category, object type, and location. The Search feature and the Inventory Browser provide visibility into your IT environment.

site

A grouping of resources generally by location. Resources are grouped by site to help with resource management and assignment.

SLA Policy

Allows storage and virtualization administrators to create customized templates for the key processes involved in the creation and use of Backup jobs. Copy types, destinations, and parameters are configured in SLA Policies, which can be used and re-used in Backup jobs.

SMTP provider

A message server that submits email notifications at the end of jobs.

snapshot

A point in time image of a resource.

Storage Protection Reports

A category of reports that help ensure your data is protected and display the status of your replication process. Use the Storage Protection Reports to view qtrees that are protected using NetApp SnapVault or NetApp SnapMirror functionality and volumes

that are protected using NetApp SnapMirror functionality. You can also view unprotected qtrees and volumes to help plan your data protection strategies.

Storage Utilization Reports

A category of reports that help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates.

Super User

The initial native user provided with the application. The Super User can perform all IBM Spectrum Copy Data Management operations.

T

tenant

A tenant is a grouping of resources and users that are administered by a tenant administrator. An IBM Spectrum Copy Data Management administrator creates tenants, assigns resources to be made available to the tenants, and creates the tenant administrator. The tenant administrator can then further control and restrict resources for users in the tenant group, as well as add additional users to the tenant through LDAP.

test mode

A Restore job mode where temporary virtual machines are created in a fenced network for such uses as development/testing, snapshot verification, and disaster recovery verification.

U

user

An IBM Spectrum Copy Data Management user, either native to IBM Spectrum Copy Data Management or imported through LDAP. The word "user" in lowercase refers generically to any type of user.

V

vault

An efficient block-level data replication technique that produces snapshot copies on secondary storage for long term scenarios such as archiving.

VMware provider

A VMware host such as an ESX/ESXi or a vCenter server.

W

workflow

A grouping of subpolicies that are sequentially joined together for creating Backup and Restore jobs. A workflow is defined visually by the storage administrator in the management interface.

Terminology Changes

The following terminology changes are implemented across the IBM Spectrum Copy Data Management user interface as of version 2.2.6.

Pre 2.2.6	2.2.6 and later
Catalog	Inventory
Copy	Backup
Hold (Job action)	Hold Schedule
Instant Access	Instant Disk Restore
Instant Virtualization	Instant VM Restore, Instant DB Restore
Pending (Job status)	Resource Active
Policy	Job, job definition
Release (Job action)	Release Schedule
Storage Workflow	SLA Policy
Use	Restore
VM Copy	VM Replication

RELATED TOPICS:

- [Documentation Roadmap](#) on page 542

Trademarks

Commonly Used Company and Product Names

Companies and products listed here may be used in the documentation:

Adobe®

PDF

DellEMC®

EMC®, EMC²®, VNX®, VNXe®

Google®

Android™, Chrome®

IBM®

FlashCopy®, FlashSystem™, IBM Spectrum™, IBM Spectrum Accelerate™, IBM Spectrum Protect Snapshot™, IBM Spectrum Virtualize™, Storwize®, XIV®

Microsoft®

Active Directory®, Excel®, Exchange, Internet Explorer®, Internet Information Services (IIS), Hyper-V®, iSCSI Initiator, SharePoint®, SourceSafe®, SQL Server®, Vista®, Visual (VSS), Windows®, Windows PowerShell®, Windows Server®, Word®

Mozilla®

Firefox®

NetApp ONTAP®

Data ONTAP®, FilerView®, FlexVol®, FlexClone®, Infinite Volume, MultiStore®, NearStore®, NOW®, OnCommand™, OnCommand Unified Manager, OSSV, RAID-DP®, Snap Creator™, SnapManager®, SnapMirror®, Snapshot™, SnapProtect®, SnapVault®, SVM, vFiler®, WAFL®

OpenLDAP™

Oracle®

Java®, JavaScript™

Pure Storage™

VMware®

ESX server, ESXi server, vCenter, vSphere, VMware Consolidated Backup, VMDK, vMotion

Additional Trademark Information

The following list contains additional trademark information:

- © 2012 Google Inc. All rights reserved. Google and Chrome are registered trademarks of Google Inc. Android is a trademark of Google Inc.
- NetApp, the NetApp logo, Go further, faster, Data ONTAP, FilerView, FlexClone, FlexVol, NearStore, RAID-DP, Snapshot, and SnapVault are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.
- Windows and SharePoint are registered trademarks of Microsoft Corporation.
- VMware is a registered trademark of VMware, Inc. Copyright © 2011 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.
- All other company and product names used herein may be the trademarks of their respective owners.

Index

A

About the Help System 543
Acronyms 598
Add a Key 128
Add a Resource Pool 106
Add a Role 109
Add a Site 72
Add Credentials to a Virtual Machine 101
Appliance Protection 560
Application Configuration Report 369
Application RPO Compliance Report 372
Application Workflow 11
Audience and Purpose 9
Authentication 14
AWS Requirements 62

B

Best Practice 155, 157, 166, 186, 253, 257, 265,
272, 320, 365, 367
Best Practices for Configuring Tenants 124
Browser Support 23

C

Catalog Browser 345
Catalog Protection 560
Catalog Summary Report 377

Character Strings 548
Clone 326, 330
Collect Logs For Troubleshooting 518
Columns 551
Configuration Report 379
Configure
 Node 75
 Provider 75
Configure Accounts 111
Configure Resource Pools 106
Configure Roles 109
Configure Scripts 152
Configure Tenants 122
Copy Data 561
Create
 Customized Report 363
 Schedule 155
Credentials 101
Customized Report 363-364

D

Dashboard Overview 67
Definitions 600
Delete
 Generated Report 367
 Schedule 159
Delete a Resource Pool 107
Delete a Role 110
Delete an Account 112

Deploy the Application 63
Deployment Checklist 19
Documentation Roadmap 542
Documentation Suite 542
Download
 Report 365
 Search Results 356

E

Edit
 Customized Report 364
 Schedule 157
Edit a Provider 98
Edit a Resource Pool 107
Edit a Role 110
Edit an Account 112
EMC Unity 70, 75
EMC VNX Catalog Data Policy Requirements 24
Encryption 15
End IV 326, 330
Export Search Results 356

F

FAQ 556
FCM 24
Fibre Channel 255, 259, 268, 275, 279, 321, 324, 328
File Analytics Reports 388
File Search 346

Files By Age Report 392
Files By Category Report 394
Files By Size Report 396
Files Usage by Owner Report 390
Filter Guidelines 548
Find and Restore a File 354
Frequently Asked Questions 556

G

Generate a Report 361
Generated Report 365
Glossary 600

H

Help 542-543
Horizontal Scale Out 534
HPE Nimble Storage 24

I

IBM Storage and IBM FCM Catalog Data Policy Requirements 24
Identification and Authentication 14
iGroup 300, 319
Install the Application 63
Installation Requirements 23
Internationalization 26
iSCSI 300, 319

J

Job History 164
Job Log Options 521
Job Logs 521
Job Monitor 164
Job Session 162, 164

K

Knowledge Base articles
 1010992 27
 1013246 27

L

LDAP
 Authentication 14, 19, 552
 Resources 95
 Server 70, 75, 86, 95, 559
 User Name Syntax 552

Legal 606
Load Balancing 248, 530-531
Load Sharing 248, 530-531
Log On to the Virtual Appliance 515
Logs For Troubleshooting 518

M

Maintenance Overview 514
Manage the Virtual Appliance 523
Manuals 542

Marketplace 527
Modify Network Settings 530
Modifying Job Log Options 521
Monitor a Job Session 164

N

Native User 14
NetApp 7-Mode 70, 75
NetApp Cluster Mode 70, 75
NetApp Knowledge Base articles
 1010992 27
 1013246 27
NetApp Protection Usage Report 417
NetApp RPO Compliance Report 420
NetApp Storage Catalog Data Policy
 Requirements 25
NetApp Storage System 70, 75
NetApp Volumes 95
Network Settings 530

O

Object Search 346
Online Help 542-543
Oracle 78
Oracle Requirements 41
OSSV Relationship Status Report 437
Overview 10

P

Pattern Matching 548
PDFs 542
Performance Considerations 248, 530-531
Ports 15
Predefined Report 360-361
Protection Compliance Reports 398
Provisioned User 14
Proxy 248, 530-531

R

Recovery Catalog 345, 357
Recovery Points Report 427
Register a Provider 75
Register a Resource 75
Registrations 606
Report Browser 361, 363-364
Reports Overview 360
Resource Browser 95
Restoring A Snapshot From A FlexGroup Volume
 To Another FlexGroup Volume 539
Return Code Reference 553
Role-Based Access Control Overview 104
RRP 326, 330
Run a Report 361

S

sAMAccountName 87

Search for Objects 346
Search Guidelines 548
Search Overview 345
Security 14
Select, Sort, and Reorder Columns 551
Set Time Zone 516
SMTP Host 70, 75
Start, Pause, and Hold a Job Session 162
Storage Protection Reports 435
Storage System 70, 75
Storage Utilization Reports 449
Subnet 323, 327
System Management Reports 376
System Requirements 23

T

Tabs 96, 349, 357
Terminology 600
Time Machine 345, 357
Time Zone 516
Trademarks 606
Trigger 155
Troubleshooting Logs 518
Truncate application logs 244

U

Unprotected Virtual Machines Report 429
Unregister a Provider 100
Upload an LDAP SSL Certificate 532

Use Data 561
User Administration and Security Management 14
User Data Security 14
User Provisioning 70, 75
Using State and Status Arguments in
 Postscripts 334

V

VADP 248, 530-531
VADP Proxy 248, 530-531
VASA 319
vCenter Host 70, 75
View a Provider 95
View File Details 352
View Object Details 351
Virtual Appliance 63, 515
Virtual Machine Privileges 545
Virtual Machines 95
VM and Storage Mapping Report 496
VMFS 242, 318
VMware Catalog 345, 357
VMware Catalog Data Policy Requirements 26
VMware Datastores Report 498
VMware Hosts 95
VMware LUNs Report 501
VMware Orphaned Datastores Report 503
VMware Orphaned LUNs Report 505
VMware RPO Compliance Report 431
VMware Server 70-71, 75, 92
VMware VM Snapshot Sprawl Report 507

VMware VM Sprawl Report 509
VMware VM Storage Report 512
VVOL 134, 243, 319

W

Wildcard Guidelines 548
Windows File System 30
Workflow 11