

# IBM Spectrum Copy Data Management<sup>™</sup> 2.2.5.0

## *User's Guide*

# Table of Contents

- Audience and Purpose** ..... 9
- IBM Spectrum Copy Data Management Overview** ..... 10
- IBM Spectrum Copy Data Management Workflow** ..... 11
- User Administration and Security Management** ..... 14
- Installation and Setup** ..... 18
  - Deployment Checklist ..... 19
  - System Requirements ..... 20
    - Virtual Machine Installation ..... 20
    - Browser Support ..... 20
    - Catalog Data Policy Requirements ..... 21
    - Copy Data and Use Data Policy Requirements ..... 24
  - Oracle Requirements ..... 31
    - Registration and Authentication ..... 31
    - Privileges ..... 31
    - Software Packages ..... 32
    - SAN Configuration ..... 33
    - Database Discovery ..... 33
    - ASM Disk Discovery ..... 33
    - Permission Considerations ..... 33
    - RMAN Considerations ..... 34
    - Name Resolution ..... 34
  - AWS Requirements ..... 35
  - Install IBM Spectrum Copy Data Management as a Virtual Appliance ..... 36
  - Start IBM Spectrum Copy Data Management ..... 38
- Dashboard** ..... 40
- Configure Sites** ..... 42
  - Sites Overview ..... 43
  - Add a Site ..... 44

Edit a Site .....	45
Delete a Site .....	46
<b>Configure Providers .....</b>	<b>47</b>
Configure Providers Overview .....	48
Register a Provider .....	50
View a Provider .....	63
Edit a Provider .....	66
Unregister a Provider .....	68
Add Credentials to a Virtual Machine .....	69
<b>Configure Role-Based Access Control .....</b>	<b>71</b>
Role-Based Access Control Overview .....	72
Configure Resource Pools .....	74
Configure Roles .....	77
Configure Accounts .....	79
VMware Admin Role-Based Access Control Configuration .....	81
Resource Pool Configuration .....	81
Role Configuration .....	82
VMware Admin Account Configuration .....	82
NetApp ONTAP Admin Role-Based Access Control Configuration .....	83
Resource Pool Configuration .....	83
Role Configuration .....	84
NetApp Admin Account Configuration .....	84
IBM Admin Role-Based Access Control Configuration .....	85
Resource Pool Configuration .....	85
Role Configuration .....	85
IBM Admin Account Configuration .....	86
<b>Configure Tenants .....</b>	<b>87</b>
<b>Best Practices for Configuring Tenants .....</b>	<b>89</b>
<b>Configure Identities .....</b>	<b>91</b>
Identities Overview .....	92
Add a Key .....	93

Add a Credential .....	96
<b>Plan .....</b>	<b>97</b>
Plan Overview .....	99
Configure Storage Workflows .....	102
Schedules .....	118
Create a Schedule .....	119
Edit a Schedule .....	121
Delete a Schedule .....	123
Catalog Data Policies .....	124
Create an Application Catalog Policy .....	125
Create a DellEMC Unity Catalog Data Policy .....	127
Create an IBM Spectrum Accelerate Catalog Data Policy .....	129
Create an IBM Spectrum Virtualize Catalog Data Policy .....	131
Create an IBM Spectrum Protect Snapshot Catalog Data Policy .....	133
Create a NetApp ONTAP Storage Catalog Data Policy .....	135
Create a NetApp ONTAP File Catalog Data Policy .....	137
Create a Pure Storage FlashArray Catalog Data Policy .....	141
Create a VMware Catalog Data Policy .....	143
Copy Data Policies .....	145
Create an Application Copy Policy .....	146
Create a DellEMC Unity Copy Data Policy .....	154
Create an IBM Spectrum Accelerate Copy Data Policy .....	157
Create an IBM Spectrum Virtualize Copy Data Policy .....	160
Create a NetApp ONTAP Copy Data Policy .....	164
Create a Pure Storage FlashArray Copy Data Policy .....	167
Create a VMware Copy Data Policy .....	170
Create VMware Copy Data Policy Proxies .....	175
Use Data Policies .....	178
Create an Application Use Policy .....	179
Create a DellEMC Unity Use Data Policy .....	186
Create an IBM Spectrum Accelerate Use Data Policy .....	191

Create an IBM Spectrum Virtualize Use Data Policy .....	196
Create a NetApp ONTAP Use Data Policy .....	201
Create a Pure Storage FlashArray Use Data Policy .....	208
Create a VMware Use Data Policy .....	213
Using State and Status Arguments in Postscripts .....	222
Create a Script Policy .....	225
Create a Report Policy .....	228
Edit a Policy .....	230
Delete a Policy .....	231
Maintenance Policy .....	232
<b>Monitor .....</b>	<b>233</b>
Monitor Overview .....	234
Start, Stop, and Hold a Job Session .....	235
Monitor a Job Session .....	237
<b>Search .....</b>	<b>239</b>
Search Overview .....	240
Search for Objects .....	241
View Object Details .....	246
View NetApp ONTAP File Details .....	247
Find and Restore a File .....	249
Download Search Results .....	251
Browse Catalog .....	252
<b>Report .....</b>	<b>254</b>
Report Overview .....	256
Run a Report .....	257
Create a Customized Report .....	259
Edit a Customized Report .....	260
Download a Report .....	261
Delete a Generated Report .....	262
Application Reports .....	263
Application Configuration Report .....	264

Application RPO Compliance Report .....	266
System Management Reports .....	270
Catalog Summary Report .....	271
Configuration Report .....	273
Policy Report .....	277
System Sizing Report .....	279
File Analytics Reports .....	281
File Usage by Owner Report .....	283
Files By Age Report .....	285
Files By Category Report .....	287
Files By Size Report .....	289
Protection Compliance Reports .....	291
DellEMC Unity RPO Compliance Report .....	293
IBM Spectrum Accelerate RPO Compliance Report .....	297
IBM Spectrum Virtualize RPO Compliance Report .....	299
NetApp ONTAP Protection Usage Report .....	303
NetApp ONTAP RPO Compliance Report .....	306
Pure Storage FlashArray RPO Compliance Report .....	310
Unprotected Virtual Machines Report .....	313
VMware RPO Compliance Report .....	315
Storage Protection Reports .....	320
NetApp ONTAP OSSV Relationship Status Report .....	322
NetApp ONTAP SnapManager Protection Status Report .....	324
NetApp ONTAP Overprotected Volumes Report .....	326
NetApp ONTAP Qtree Protection Status Report .....	328
NetApp ONTAP Underprotected Volumes Report .....	330
NetApp ONTAP Volume Protection Status Report .....	332
NetApp ONTAP Transition Dependency Report .....	334
Storage Utilization Reports .....	336
DellEMC Unity File Systems Report .....	339
DellEMC Unity LUNs Report .....	341

DellEMC Unity Pools Report .....	344
IBM Spectrum Accelerate Pools Report .....	346
IBM Spectrum Accelerate Volumes Report .....	348
IBM Spectrum Virtualize Consistency Groups Report .....	350
IBM Spectrum Virtualize Pools Report .....	352
IBM Spectrum Virtualize Volumes Report .....	355
NetApp ONTAP Aggregates Report .....	358
NetApp ONTAP LUNs Report .....	361
NetApp ONTAP Orphaned LUNs Report .....	363
NetApp ONTAP Quotas Report .....	365
NetApp ONTAP Snapshots Report .....	367
NetApp ONTAP Volumes Report .....	369
Pure Storage FlashArray Volumes Report .....	371
Storage Capacity Report .....	374
VMware Datastores Report .....	375
VMware LUNs Report .....	377
VMware Orphaned Datastores Report .....	379
VMware Orphaned LUNs Report .....	381
VMware VM Snapshot Sprawl Report .....	383
VMware VM Sprawl Report .....	384
<b>Maintenance .....</b>	<b>387</b>
Maintenance Overview .....	388
Log On to the Virtual Appliance .....	389
Set Time Zone .....	390
Collect Logs For Troubleshooting .....	391
Manage the Virtual Appliance .....	393
Update IBM Spectrum Copy Data Management .....	394
Install the Marketplace RPM .....	396
Upload an SSL Certificate .....	397
<b>Documentation and Support .....</b>	<b>399</b>
Documentation Roadmap .....	400

---

About the Help System .....	401
<b>Reference Topics .....</b>	<b>403</b>
Search and Filter Guidelines .....	404
Select, Sort, and Reorder Columns .....	407
LDAP User Name Syntax .....	408
Return Code Reference .....	409
<b>Frequently Asked Questions .....</b>	<b>412</b>
<b>Acronyms .....</b>	<b>419</b>
<b>Terminology .....</b>	<b>421</b>
<b>Trademarks .....</b>	<b>426</b>
<b>Index .....</b>	<b>428</b>



# 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, and schedule jobs.

# IBM Spectrum Copy Data Management Overview

IBM Spectrum Copy Data Management is an application that delivers visibility, insight, and control into your IT environment. It provides an enterprise-wide view of IT objects and files, so you can discover, use, copy, search, and report on data, giving you greater understanding of your organization. IBM Spectrum Copy Data Management catalogs files and other IT objects making them easily searchable, and provides the ability to distinguish where data copies are located in your IT environment and leverage the most appropriate copy for any given use case. It catalogs information on both primary and secondary storage, across all copy instances and replicas. IBM Spectrum Copy Data Management rapidly searches through storage providers on a network to identify objects such as files, directories, volumes, and qtrees. Objects can be found across all tiers of information in a single search, giving you comprehensive visibility of your objects, their copies, and their locations.

IBM Spectrum Copy Data Management utilizes automated Copy Data Management workflows for replicating and intelligently reusing snapshots, vaults, and mirrors. Copy Data and Use Data policies offer control over testing and cloning use cases, instant recovery, and full disaster recovery. Through Copy Data and Use Data policies, you can:

- Copy data from a variety of storage providers to multiple locations.
- Reuse and recover resources from snapshots, vaults, and mirrors.
- Support use cases for automated data protection, recovery, DevOps, Dev/Test, data and database validation with data masking, through the use of automated Instant Access, Instant Virtualization, volume, and file restore functionalities.

IBM Spectrum Copy Data Management also provides a wide range of reports on cataloged information, including resource utilization, protection, and file-level statistics. It lets you quickly validate whether your critical data is protected, as well as alerts you of over-protected data.

# 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 Catalog by registering them.

1. Click the **Configure**  tab and the **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 a Catalog Data Policy***

A Catalog Data policy provides the framework to collect and catalog information about objects on a registered provider.




1. Click the **Plan**  tab and the **Policies**  view.
2. Click **New**  and then select a catalog data policy.
3. Select one or more resources to catalog from the list of available providers.
4. Select the options for your policy. 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 policy and save the policy as a job.



## Run a Job

A job that is based on a Catalog Data policy discovers object information, catalogs it, and populates the IBM Spectrum Copy Data Management database.

1. Click the **Monitor**  tab and the **Jobs**  view.
2. Select the job to run by clicking in the row containing the job name.
3. Click **Start** . The job session runs.





## 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.

## Copy Data

Create copies of your data. The RPO and copy data parameters are defined in a Storage Workflow, which is then applied to the Copy Data policy along with a specified activation time to meet your copy data criteria.

1. Click the **Plan**  tab and the **Policies**  view.
2. Click **New**  and then select a Copy Data policy.

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

## ***Use Data***

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

1. Click the **Plan** 📅 tab and the **Policies** 📄 view.
2. Click **New** ➕ and then select a Use Data policy and workflow template.
3. Click the Workflow tab and select providers to reuse or recover, as well as destinations.
4. Complete the policy definition including notification and other options. Save the policy.
5. Click **Run** ▶ . The selected data is made available for use in accordance with the defined policy parameters.

### **RELATED TOPICS:**

- [Start IBM Spectrum Copy Data Management](#) on page **38**
- [Configure Providers Overview](#) on page **48**
- [Plan Overview](#) on page **99**
- [Start, Stop, and Hold a Job Session](#) on page **235**
- [Search for Objects](#) on page **241**
- [Report Overview](#) on page **256**

# User Administration and Security Management

IBM Spectrum Copy Data Management provides users the opportunity to rapidly locate files and objects on DellEMC, IBM, and NetApp storage 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 and/or NetApp 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, 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, and NetApp 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	Version	Comment
22	ssh	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	VMware and NetApp ONTAP	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
5986	WinRM	Windows Remote Management	Service used by IBM Spectrum Copy Data Management



Port	Service	Version	Comment
8443	ssl/http	Apache Tomcat/Coyote JSP engine 1.1	Port open within the firewall
8761	Discovery Server	1.0	Service used by IBM Spectrum Copy Data Management. Locates registered micro services.
27017	MongoDB	MongoDB mongod	Service used by IBM Spectrum Copy Data Management
27018	MongoDB	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 50
- [Role-Based Access Control Overview](#) on page 72

---

# Installation and Setup

[Deployment Checklist](#) on page **19**

[System Requirements](#) on page **20**

[Oracle Requirements](#) on page **31**

[AWS Requirements](#) on page **35**

[Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page **36**

[Start IBM Spectrum Copy Data Management](#) on page **38**

[Dashboard](#) on page **40**

## 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.

### 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.	<a href="#">Install IBM Spectrum Copy Data Management as a Virtual Appliance</a> on page <b>36</b>	
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.	<a href="#">Install IBM Spectrum Copy Data Management as a Virtual Appliance</a> on page <b>36</b>	
3	Log on to the IBM Spectrum Copy Data Management virtual appliance to set your local time zone. Time zone is used for policy scheduling.	<a href="#">Set Time Zone</a> on page <b>390</b>	
4	Launch IBM Spectrum Copy Data Management to set a new Super User password.	<a href="#">Start IBM Spectrum Copy Data Management</a> on page <b>38</b>	
5	Configure LDAP Authentication and establish a secure connection to the LDAP server.	<a href="#">Register a Provider</a> on page <b>50</b>	
6	Create accounts through role-based access control.	<a href="#">Role-Based Access Control Overview</a> on page <b>72</b>	

### Post-Deployment Checklist

Step	Action	Related Topic	✓
•	Launch IBM Spectrum Copy Data Management to begin using the application and its features.	<a href="#">Start IBM Spectrum Copy Data Management</a> on page <b>38</b>	
•	Update IBM Spectrum Copy Data Management to keep it current with new features, enhancements, and upgrades.	Internet Update or CD Update	
•	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.	<a href="#">Manage the Virtual Appliance</a> on page <b>393</b>	

# System Requirements

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

## 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 5.1, 5.5, or 6.0
- 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
- 32 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 policies, jobs, events, and logs
- 250 GB for Catalog data

## Browser Support

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

IBM Spectrum Copy Data Management 2.5 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.

**Note:** Pop-up windows must be enabled in your browser to access the Help system and some IBM Spectrum Copy Data Management operations.

## Catalog Data Policy Requirements

The following requirements apply to Catalog Data policies.

### Application Catalog Policy Requirements for Oracle Databases

Oracle 11gR2 and Oracle 12c on a physical or virtual machine running the following operating systems:

- Red Hat Enterprise Linux/CentOS 6.5 and later
- Red Hat Enterprise Linux/CentOS 7.0 and later
- Oracle Enterprise Linux 6.0 and above, 7.0 and later
- SUSE Linux Enterprise Server 11.0 Service Pack 4 and later, 12.0 and later

Oracle standalone, RAC, and ASM configurations are supported. 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.

**Note:** Oracle support for virtual machines requires iSCSI LUNs connected directly to the guest system. Oracle configured with RDM LUNs or VMDKs in VM is not supported.

**Note:** 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.

### Supported Storage Systems for Oracle Databases

- IBM storage systems running IBM Spectrum™ Virtualize Software version 7.3 and later, including IBM SAN Volume Controller, IBM Storwize, and IBM FlashSystem V9000 systems. IBM providers must be registered by an IBM user with administrator-level privileges. For more information, see [Oracle Requirements](#) on page 31.
- Pure Storage FlashArray//m and FlashArray 4xx series

### Application Catalog Policy Requirements for Microsoft SQL Server

SQL Server 2012, SQL Server 2014, and SQL Server 2016 standalone and AlwaysON, installed to a virtual machine using a VMDK configuration, running Microsoft Windows 2012 R2

**Note:** It is highly recommended to install the latest SQL Server patches and updates in your environment.

### Supported Storage Systems for Microsoft SQL Server

IBM Spectrum Copy Data Management supports VADP-based VM Copies to any of the following storage system targets:

- IBM storage systems running IBM Spectrum™ Virtualize Software version 7.3 and later,

including IBM SAN Volume Controller, IBM Storwize, and IBM FlashSystem V9000 systems

- IBM storage system running IBM Spectrum Accelerate version 11.5.3 and later, including IBM FlashSystem A9000/A9000R and IBM XIV storage systems
- Pure Storage FlashArray//m and FlashArray-4xx series
- NetApp Storage system running Data ONTAP software

Note the following Microsoft environmental requirements:

- Windows Remote Shell (WinRM) must be enabled
- 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.
- The SQL user must enable the **public** and **sysadmin** SQL permissions.
- The user identity must have sufficient rights to install and start the ECX Tools Service on the virtual machine node. This includes "on as a service" rights.
- The clock skew between the Domain Controller and the IBM Spectrum Copy Data Management appliance should not be more than 5 minutes
- 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

### **DellEMC Unity Catalog Data Policy 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

### **IBM Storage and IBM FCM Catalog Data Policy Requirements**

- IBM storage systems running IBM Spectrum™ Virtualize Software version 7.3 and later, including IBM SAN Volume Controller, IBM Storwize, and IBM FlashSystem V9000 systems
- 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 Storage Catalog Data Policy 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

**Note:** Clustered Data ONTAP providers must be registered with a cluster administrator account.

**Note:** For IBM Spectrum Copy Data Management 2.2 and later, 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 Catalog Data Policy Requirements

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

- 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 8.2p3 and later
- Clustered Data ONTAP 8.3 and later

**Note:** Clustered Data ONTAP providers must be registered with a cluster administrator account.

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

- **create\_unicode** and **convert\_unicode**

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.

**Note:** 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 Catalog Data Policy Requirements

- Pure Storage FlashArray//m and FlashArray-4xx series
- Pure Storage REST API version 1.6 and later

### VMware Catalog Data Policy Requirements

- vSphere 5.1 and later
- vSphere 5.5 and later
- vSphere 6.0 and later

### Copy Data and Use Data Policy Requirements

The following requirements apply to Copy Data and Use Data policies.

### Application Copy and Use Policy Requirements for Oracle Databases

Oracle 11gR2 and Oracle 12c on a physical or virtual machine running the following operating systems:

- Red Hat Enterprise Linux/CentOS 6.5 and later
- Red Hat Enterprise Linux/CentOS 7.0 and later
- Oracle Enterprise Linux 6.0 and above, 7.0 and later
- SUSE Linux Enterprise Server 11.0 Service Pack 4 and later, 12.0 and later

Oracle standalone, RAC, and ASM configurations are supported. 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.



**Note:** Oracle support for virtual machines requires iSCSI LUNs connected directly to the guest system. Oracle configured with RDM LUNs or VMDKs in VM is not supported.

**Note:** 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.

### Supported Storage Systems for Oracle Databases

- IBM storage systems running IBM Spectrum™ Virtualize Software version 7.3 and later, including IBM SAN Volume Controller, IBM Storwize, and IBM FlashSystem V9000 systems. IBM providers must be registered by an IBM user with administrator-level privileges. For more information, see [Oracle Requirements](#) on page 31.
- Pure Storage FlashArray//m series and FA-4xx series

### Application Copy and Use Policy Requirements for Microsoft SQL Server

SQL Server 2012, SQL Server 2014, and SQL Server 2016 standalone and AlwaysON, installed to a virtual machine using a VMDK configuration, running Microsoft Windows 2012 R2

**Note:** It is highly recommended to install the latest SQL Server patches and updates in your environment.

### Supported Storage Systems for Microsoft SQL Server

IBM Spectrum Copy Data Management supports VADP-based VM Copies to any of the following storage system targets:

- IBM storage systems running IBM Spectrum™ Virtualize Software version 7.3 and later, including IBM SAN Volume Controller, IBM Storwize, and IBM FlashSystem V9000 systems
- IBM storage system running IBM Spectrum Accelerate version 11.5.3 and later, including IBM FlashSystem A9000/A9000R and IBM XIV storage systems
- Pure Storage FlashArray//m and FA-4xx series
- NetApp Storage system running Data ONTAP software

**Note:** SQL servers residing on any storage can also be protected to supported storage systems through VADP-based VM Copies.

Note the following Microsoft environmental requirements:

- Windows Remote Shell (WinRM) must be enabled
- 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.
- The SQL user must enable the **public** and **sysadmin** SQL permissions.

- The user identity must have sufficient rights to install and start the ECX Tools Service on the virtual machine node. This includes "Log on as a service" rights.
- The clock skew between the Domain Controller and the IBM Spectrum Copy Data Management appliance should not be more than 5 minutes
- 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

### DelEMC Unity Copy and Use Data Policy 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

### IBM Storage Copy and Use Data Policy Requirements

- IBM storage systems running IBM Spectrum™ Virtualize Software version 7.3 and later, including IBM SAN Volume Controller, IBM Storwize, and IBM FlashSystem V9000 systems
- IBM storage systems running IBM Spectrum Accelerate version 11.5.3 and later, including FlashSystem A9000/A9000R and IBM XIV storage systems

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

### NetApp ONTAP Copy Data and Use Data Policy Requirements

- Clustered Data ONTAP 8.1 and later operating in 7-Mode
- Clustered Data ONTAP 8.2 and later
- Clustered Data ONTAP 8.3 and later

**Note:** Clustered Data ONTAP providers must be registered with a cluster administrator account.

**Note:** For IBM Spectrum Copy Data Management 2.2 and later, 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.

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](#).

### Pure Storage Copy Data and Use Data Policy Requirements

- Pure Storage FlashArray//m and FA-4xx series
- Pure Storage REST API version 1.6 and later

### VMware Copy Data and Use Data Policy Requirements

- vSphere 5.1 and later
- vSphere 5.5 and later
- vSphere 6.0 and later

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 Access and Instant Virtualization 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.

### iSCSI Requirements

Both Instant Access and Instant Virtualization Use Data policies require iSCSI for VMFS, VMware's cluster file system.

### Virtual Machine Privileges

The following virtual machine privileges are required for Copy Data and Use Data policies if a virtual machine is configured as a provider in IBM Spectrum Copy Data Management with credentials lower than an administrator. If administrator credentials are used when adding the provider to IBM Spectrum Copy Data Management, these privileges are automatically enabled.

Operations

<ul style="list-style-type: none"><li>• Folder.Create</li><li>• VirtualMachine.Inventory.Create</li><li>• Host.Config.Storage</li><li>• System.Read</li><li>• Datastore.Rename</li><li>• VirtualMachine.Inventory.Delete</li><li>• VirtualMachine.Inventory.Unregister</li></ul>	<ul style="list-style-type: none"><li>• Resource.HotMigrate</li><li>• VirtualMachine.State.RemoveSnapshot</li><li>• VirtualMachine.State.RevertToSnapshot</li><li>• VirtualMachine.Interact.PowerOn</li><li>• VirtualMachine.Interact.PowerOff</li><li>• System.View</li></ul>

Reconfiguring the virtual machine may require the following privileges:

### Virtual Machine Reconfigure

<ul style="list-style-type: none"> <li>VirtualMachine.Interact.DeviceConnection if changing the runtime connection state of a device as embodied by the Connectable property.</li> <li>VirtualMachine.Interact.SetCDMedia if changing the backing of a CD-ROM device</li> <li>VirtualMachine.Interact.SetFloppyMedia if changing the backing of a floppy device</li> <li>VirtualMachine.Config.Rename if renaming the virtual machine</li> <li>VirtualMachine.Config.Annotation if setting an annotation value</li> <li>VirtualMachine.Config.AddExistingDisk if adding a virtual disk device that is backed by an existing virtual disk file</li> <li>VirtualMachine.Config.AddNewDisk if adding a virtual disk device for which the backing virtual disk file is to be created</li> <li>VirtualMachine.Config.RemoveDisk if removing a virtual disk device that refers to a virtual disk file</li> <li>VirtualMachine.Config.CPUCount if changing the number of CPUs</li> <li>VirtualMachine.Config.Memory if changing the amount of memory</li> <li>VirtualMachine.Config.RawDevice if adding, removing, or editing a raw device mapping (RDM) or SCSI passthrough device</li> <li>VirtualMachine.Config.AddRemoveDevice if adding or removing any device other than disk, raw, or USB device</li> <li>VirtualMachine.Config.EditDevice if changing the settings of any device</li> <li>VirtualMachine.Config.Settings if changing any</li> </ul>	<ul style="list-style-type: none"> <li>VirtualMachine.Config.SwapPlacement if changing swapPlacement</li> <li>VirtualMachine.Config.HostUSBDevice if adding, removing, or editing a VirtualUSB device backed by the host USB device</li> <li>VirtualMachine.Config.DiskExtend if extending an existing VirtualDisk device</li> <li>VirtualMachine.Config.ChangeTracking if enabling or disabling changed block tracking for the virtual machine's disks</li> <li>VirtualMachine.Config.MksControl if toggling display connection limits or the guest auto-lock feature</li> <li>DVSwitch.CanUse if connecting a VirtualEthernetAdapter to a port in a DistributedVirtualSwitch.</li> <li>DVPortgroup.CanUse if connecting a VirtualEthernetAdapter to a DistributedVirtualPortgroup.</li> <li>Creating a virtual machine may require the following privileges:                         <ul style="list-style-type: none"> <li>VirtualMachine.Config.RawDevice if adding a raw device</li> <li>VirtualMachine.Config.AddExistingDisk if adding a VirtualDisk and the fileOperation is unset</li> <li>VirtualMachine.Config.AddNewDisk if adding a VirtualDisk and the fileOperation is set</li> <li>VirtualMachine.Config.HostUSBDevice if adding a VirtualUSB device backed by the host USB device</li> </ul> </li> </ul>

<p>basic settings such as those in ToolsConfigInfo, FlagInfo, or DefaultPowerOpInfo</p> <ul style="list-style-type: none"><li>• VirtualMachine.Config.Resource if changing resource allocations, affinities, or setting network traffic shaping or virtual disk shares</li><li>• VirtualMachine.Config.AdvancedConfig if changing values in extraConfig</li></ul>	<p>In addition, this operation may require the following privileges:</p> <p>Datastore.AllocateSpace on any datastore where virtual disks will be created or extended</p> <p>Network.Assign on any network the virtual machine will be connected to</p>

**RELATED TOPICS:**

- [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page **36**
- [Oracle Requirements](#) on page **31**
- [AWS Requirements](#) on page **35**
- [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](#)

# Oracle Requirements

Review the following requirements and pre-requisites for registering an Oracle provider in IBM Spectrum Copy Data Management.

## Registration and Authentication

Each individual Oracle server should be registered as a provider in IBM Spectrum Copy Data Management. Like any other provider, it can be registered by name or IP address. Ensure the name is resolvable by IBM Spectrum Copy Data Management.

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). Ensure the server has the SSH service running and configure any firewalls to allow IBM Spectrum Copy Data Management the ability to connect to the Oracle server through SSH.

IBM Spectrum Copy Data Management can connect to the Oracle server as a local operating system user using a password or an SSH key. To use a password, select or create a credential of type "Local". To use a key, enter a username and select or create an SSH key. See [Identities Overview](#) on page 92.

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.

## Privileges

IBM Spectrum Copy Data Management connects to the Oracle server 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 requires the following privileges:

- **Root privileges for the operating system** - Required for tasks such as discovering storage information, rescanning iSCSI sessions, and mounting and unmounting disks.
- **Permissions to read the Oracle inventory** - Required to discover and collect information about Oracle homes and databases.
- **SYSDBA privileges for database instances** - Required to perform several database tasks such as querying instance details, starting and ending hot backup mode and performing RMAN cataloging during backups, as well as starting and stopping instances during restore.
- **SYSASM privileges for ASM instances** - Required to perform several storage tasks such as querying ASM disk information as well as renaming, mounting, and unmounting diskgroups during restore.

To ensure that the local user has each of the necessary privileges:

1. Configure `sudo` to allow the user to run commands without a password and without a `tty`.
2. Add the user to the operating system group that owns the Oracle inventory (typically `oinstall`).
3. Add the user to the `OSDBA` operating system group for each Oracle Home (typically `dba`).
4. Add the user to the `OSASM` operating system group for the Grid Home (typically `asmadmin`). This step is not required if ASM is not in use on the server.

For example, to create a new operating system user named `ecxagent` for IBM Spectrum Copy Data Management to log in to the Oracle server and grant it the necessary privileges:

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

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

- Add the user to the Oracle installation and `OSDBA` group, noting that the command may vary depending on your version or distribution of Linux: `usermod -a -G oinstall,dba ecxagent`
- If ASM is in use, also add the user to the `OSASM` group: `usermod -a -G asmadmin ecxagent`
- Place the following lines in your `sudoers` configuration file, typically `/etc/sudoers`:

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

## Software Packages

If the Oracle server runs RHEL/OEL/CentOS 6.x, ensure that 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` instead of `util-linux-ng`.

If Oracle data or logs reside on LVM volumes, ensure the LVM version is 2.0.2.118 or later. Run `lvm version` to check the version, and if necessary, run `yum update lvm2` to update the package.

Note that an automated disk mount of an IBM Spectrum Copy Data Management copy on LVM storage must not have the original disk for that copy still present and mounted while the `lvm2-lvmetad` service is active on the system. If Oracle data resides on LVM volumes, you must stop and disable the `lvm2-lvmetad` service before running Application Copy or Use jobs. To disable the `lvm2-lvmetad`, 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
```

## SAN Configuration

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

## Database Discovery

IBM Spectrum Copy Data Management discovers Oracle databases by reading the file `/etc/oratab` on each Oracle server. When you create databases using the Oracle Database Configuration Assistant (dbca), it automatically creates appropriate entries in `/etc/oratab`. However, if a database was created manually or through custom scripts, you must ensure that you also create the necessary `oratab` entries. Examine the existing `/etc/oratab` file on your system for sample entries.

Note that while each `oratab` entry begins with the `ORACLE_SID` for standalone databases, it must instead begin with the `DB_UNIQUE_NAME` for RAC databases. For example, if the global database name is `MYTESTDB`, the local `SIDs` on the RAC nodes may be named `MYTESTDB1`, `MYTESTDB2`, etc. However, the `oratab` entry must contain the global name `MYTESTDB`. For example:

```
MYTESTDB:/u01/app/oracle/product/12.1.0/dbhome_1:N
```

Refer to Oracle documentation for details on `oratab` configuration for standalone and RAC databases.

## ASM Disk Discovery

In order to mount clones/copies of ASM disks, IBM Spectrum Copy Data Management automatically creates `udev` rules on the Oracle server to change the ownership of the newly mapped disks. The `udev` rules ensure that the disks are owned by the owner of the Grid Home so that the ASM instance can discover and mount them. They also ensure that for each ASM disk mapped by IBM Spectrum Copy Data Management, a symbolic link of the form `/dev/ecx-asmdisk/<diskuuid>` is created and points to the parent disk.

Depending on the existing settings of the `ASM_DISKSTRING` parameter, ASM may not always be able to discover the cloned/copied disks. To ensure that ASM is able to discover the disks mapped by IBM Spectrum Copy Data Management, modify the existing `ASM_DISKSTRING` parameter and append the following value:  
`/dev/ecx-asmdisk/*`.

**Note:** 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, and then append the value above.

Refer to Oracle documentation for details about retrieving and modifying the `ASM_DISKSTRING` parameter.

## Permission 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.

## RMAN Considerations

IBM Spectrum Copy Data Management does not manage retention policies for RMAN catalogs or archived log locations including any additional log destinations that may be enabled as part of an IBM Spectrum Copy Data Management copy policy. Only the retention policies specified by the database administrator will apply.

## Name Resolution

For Oracle RAC environments, there may be rare cases where IBM Spectrum Copy Data Management catalogs a particular cluster node but is unable to discover the addresses of its peer nodes. In such cases, IBM Spectrum Copy Data Management attempts to directly resolve the peer nodes using their short DNS names. If IBM Spectrum Copy Data Management logs name resolution errors when running catalog jobs, you may need to specify custom DNS suffixes by running a pre-installed script through an IBM Spectrum Copy Data Management Script policy.

1. Click **Plan**  tab. On the View pane select **Policies** , then click New.
2. Select **Script**  in the System column. The Script policy editor opens.
3. In the **Enter a new command** field, enter the following:

```
sudo /opt/ECX/tools/scripts/update_resolv.py [DNS address to search]
```

Note: Separate multiple addresses with a space.

For example, `sudo /opt/ECX/tools/scripts/update_resolv.py abc.site.us xyz.site.us`

4. Press **Enter** to add the command to the script policy.
5. In the **Schedule** tab, schedule a time to run the policy, or select **Start job now** to run the job once the policy creation is complete.
6. Click **Finish** to save the policy.

### RELATED TOPICS:

- [Register a Provider](#) on page **50**
- [Identities Overview](#) on page **92**
- [System Requirements](#) on page **20**

## 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](https://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](https://docs.aws.amazon.com/storagegateway/latest/userguide/managing-cache-common.html)

[docs.aws.amazon.com/storagegateway/latest/userguide/managing-upload-buffer-common.html](https://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 policy is created. Note that 32 volumes are supported per storage gateway.

### RELATED TOPICS:

- [Register a Provider](#) on page **50**
- [Identities Overview](#) on page **92**
- [System Requirements](#) on page **20**

## 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.
- 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**.
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 on the IBM Spectrum Copy Data Management virtual appliance through vSphere Client. See [Set Time Zone](#) on page 390.
- 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 38.


#### RELATED TOPICS:

- [Start IBM Spectrum Copy Data Management](#) on page 38
- [Set Time Zone](#) on page 390

# 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 36.
- 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.
- By default, the Help system accessed from the management interface is online. To ensure you are seeing the latest Help system, clear your browser cache before selecting Help. To use a local Help system that does not require an Internet connection click the arrow next to **Help**  to open the About IBM Spectrum Copy Data Management window. Select **Use local help system**. The local version might not contain the latest updates.

## 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 72.
- Add storage systems and virtual machine resources to the IBM Spectrum Copy Data Management database. See [Register a Provider](#) on page 50 and [Plan Overview](#) on page 99.
- Search or browse for objects that match certain criteria. See [Search for Objects](#) on page 241 and [Browse Catalog](#) on page 252.

- Generate reports with predefined or customized parameters. See [Report Overview](#) on page **256**.

**RELATED TOPICS:**




- [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page **36**
- [System Requirements](#) on page **20**

# 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

### Catalog 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 Catalog 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 file policies that are no longer in use and delete them, then run the Maintenance policy 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:

- [Configure Providers Overview](#) on page 48
- [Role-Based Access Control Overview](#) on page 72
- [Plan Overview](#) on page 99
- [Monitor a Job Session](#) on page 237
- [Search Overview](#) on page 240
- [Report Overview](#) on page 256

# Configure Sites

[Sites Overview](#) on page **43**

[Add a Site](#) on page **44**

[Edit a Site](#) on page **45**

[Delete a Site](#) on page **46**

## Sites Overview

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 policies.

Sites are assigned when registering providers. When creating Copy Data and Use Data policies, sites clearly identify where your data is replicated by location.

### RELATED TOPICS:

- [Add a Site](#) on page **44**
- [Edit a Site](#) on page **45**
- [Delete a Site](#) on page **46**

## 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 Association to view the resources that are associated with the site.

### BEFORE YOU BEGIN:

- Review the properties and location of your current providers. See [Configure Providers Overview](#) on page 48 for a list of supported providers.

### To add a site:

1. Click the **Configure**  tab. On the Views pane, select **Sites** . 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 50 and [Edit a Provider](#) on page 66.


### RELATED TOPICS:

- [Sites Overview](#) on page 43
- [Edit a Site](#) on page 45
- [Delete a Site](#) on page 46




## Edit a Site

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

### BEFORE YOU BEGIN:

- Review the properties of your current sites on the **Sites**  pane.

### To edit a site:

1. Click the **Configure**  tab. On the Views pane, select **Sites** . 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 50 and [Edit a Provider](#) on page 66.

### RELATED TOPICS:




- [Sites Overview](#) on page 43
- [Add a Site](#) on page 44
- [Delete a Site](#) on page 46

## Delete a Site

Delete a site when it becomes obsolete.

A site cannot be deleted if it is assigned to a provider. Re-assign your providers to different sites before deleting.

### To delete a site:

1. Click the **Configure**  tab. On the Views pane, select **Sites** . 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 **50** and [Edit a Provider](#) on page **66**.

### RELATED TOPICS:

- [Sites Overview](#) on page **43**
- [Add a Site](#) on page **44**
- [Edit a Site](#) on page **45**

# Configure Providers

[Configure Providers Overview](#) on page **48**

[Register a Provider](#) on page **50**


[View a Provider](#) on page **63**

[Edit a Provider](#) on page **66**

[Unregister a Provider](#) on page **68**

[Add Credentials to a Virtual Machine](#) on page **69**

## Configure Providers Overview

Add providers to IBM Spectrum Copy Data Management and define user access roles on the **Configure**  tab.

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.
- 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, including IBM SAN Volume Controller, IBM Storwize, and IBM FlashSystem V9000 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.
- 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 policies 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 **50**
- [View a Provider](#) on page **63**



- [Edit a Provider](#) on page **66**
- [Unregister a Provider](#) on page **68**

## 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.
- 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, including IBM SAN Volume Controller, IBM Storwize, and IBM FlashSystem V9000 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.
- 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 **Catalog provider resources after registration** option is selected, IBM Spectrum Copy Data Management creates a high-level catalog policy and automatically catalogs the objects on the provider.

**Note:** For IBM Spectrum Copy Data Management 2.2 and later, 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 44.

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

The following describes the procedures for registering a provider.

### To register an LDAP provider:

1. Click the **Configure**  tab. On the Views pane, select **Providers** . The Provider Browser opens.
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.

#### Bind DN

The name used for authenticating the connection to the LDAP server. IBM Spectrum Copy Data Management supports simple bind.

#### Password

The password associated with the Bind Distinguished Name.

#### 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 **Providers** . The Provider Browser opens.
2. In the Provider Browser pane, select **NetApp ONTAP** .
3. Right-click **NetApp ONTAP** . Then click **Register** . The Register NetApp ONTAP dialog opens.
4. Populate the fields in the dialog:

#### Site

A user-defined provider location, created in the Sites 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.

#### Username

The name used to access the provider.

#### Password

The password associated with the user name.

#### Comment

Optional provider description.

### Catalog provider resources after registration





If selected, IBM Spectrum Copy Data Management creates a high-level catalog policy and automatically catalogs the objects on the provider. Note that the catalog policy 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 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 **Providers** . The Provider Browser opens.
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, username, and password. 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 **Catalog provider resources after registration** is selected, IBM Spectrum Copy Data Management creates a high-level catalog policy 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 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.
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.

**Username**

The name used to access the provider.

**Password**

The password associated with the user name.

**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.
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 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.

**Username**

The name used to access the provider.

**Password**

The password associated with the user name.

**Comment**

Optional provider description.

**Catalog provider resources after registration**

If selected, IBM Spectrum Copy Data Management creates a high-level catalog policy and automatically catalogs the objects on the provider. Note that the catalog policy 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 69.

**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 **Providers** . The Provider Browser opens.
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, username, and password. 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 **Catalog provider resources after registration** is selected, IBM Spectrum Copy Data Management creates a high-level catalog policy 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 69.

### **To register an IBM Spectrum Protect Snapshot or IBM Spectrum Virtualize provider:**

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

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

#### **Site**

A user-defined provider location, created in the Sites 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.

#### **Username**

The name used to access the provider.

#### **Password**

The password associated with the user name.

#### **Comment**



Optional provider description.

### Catalog provider resources after registration

If selected, IBM Spectrum Copy Data Management creates a high-level catalog policy and automatically catalogs the objects on the provider. Note that the catalog policy 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.

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 Accelerate provider:

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

#### Site

A user-defined provider location, created in the Sites 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.

### Catalog provider resources after registration

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

#### Credentials






Select or create your IBM Spectrum Accelerate credentials. See [Identities Overview](#) on page 92.

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 a DellEMC Unity provider:**

1. Click the **Configure**  tab. On the Views pane, select **Providers** . The Provider Browser opens.
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 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.

#### **Username**

The name used to access the provider.

#### **Password**

The password associated with the user name.

#### **Comment**

Optional provider description.

#### **Catalog provider resources after registration**





If selected, IBM Spectrum Copy Data Management creates a high-level catalog policy and automatically catalogs the objects on the provider. Note that the catalog policy 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 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 **Providers** . The Provider Browser opens.
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, username, and password. 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 **Catalog provider resources after registration** is selected, IBM Spectrum Copy Data Management creates a high-level catalog policy 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 Oracle application provider:**

1. Click the **Configure**  tab.
2. In the Provider Browser pane, select **Application**.
3. Right-click **Application**. 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 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).

#### **Catalog provider resources after registration**

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

#### **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 92.

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.

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 SQL application provider:**

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

1. Click the **Configure**  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 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.

#### **vCenter**

The vCenter location of the SQL application server.



#### **Credentials**



Select or create your SQL credentials. See [Identities Overview](#) on page 92.

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 Amazon Web Services (AWS) provider:**

1. Click the **Configure**  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 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 92.

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 provider:**

1. Click the **Configure**  tab. On the Views pane, select **Providers** . The Provider Browser opens.
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 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 **92**.

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.

#### 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 **74**.
- Add credentials to virtual machines in a VMware environment. See [Add Credentials to a Virtual Machine](#) on page **69**.

#### RELATED TOPICS:

- [Configure Providers Overview](#) on page **48**
- [View a Provider](#) on page **63**
- [Edit a Provider](#) on page **66**
- [Unregister a Provider](#) on page **68**
- [Sites Overview](#) on page **43**
- [Upload an SSL Certificate](#) on page **397**
- [LDAP User Name Syntax](#) on page **408**



## 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 50.

### To view a list of providers:

1. Click the **Configure**  tab. On the Views pane, select **Providers** . The Provider Browser opens.
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.

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:**

- Catalog storage providers that have not been recently cataloged. See [Plan Overview](#) on page 99.

### **RELATED TOPICS:**



- [Configure Providers Overview](#) on page **48**
- [Register a Provider](#) on page **50**
- [Edit a Provider](#) on page **66**
- [Unregister a Provider](#) on page **68**




## Edit a Provider

Revise the properties of a provider as needed.

### BEFORE YOU BEGIN:

- Review the properties of your current providers. See [Configure Providers Overview](#) on page 48 for a list of supported providers.

### To edit the properties of a provider:

1. Click the **Configure**  tab. On the Views pane, select **Providers** . The Provider Browser opens.
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 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.

#### Username

The name used to access the provider.

#### Password

The password associated with the user name.

#### Comment

Optional provider description.

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

### NEXT STEPS:

- If the storage provider you edited has not recently been cataloged, catalog it. See [Plan Overview](#) on page **99**.

**RELATED TOPICS:**

- [Configure Providers Overview](#) on page **48**
- [Register a Provider](#) on page **50**
- [View a Provider](#) on page **63**
- [Unregister a Provider](#) on page **68**

## 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 policy.

**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 policy.

### BEFORE YOU BEGIN:

- Review the properties of the provider to determine if you want to unregister it. See [Edit a Provider](#) on page 66.

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

### To unregister a provider:

1. Click the **Configure**  tab. On the Views pane, select **Providers** . The Provider Browser opens.
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:

- [Configure Providers Overview](#) on page 48
- [Register a Provider](#) on page 50
- [View a Provider](#) on page 63
- [Edit a Provider](#) on page 66




## 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 Copy Data policy. 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 50 and [Create a VMware Catalog Data Policy](#) on page 143.


### To add credentials for a single virtual machine:

1. Click the **Configure**  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.
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 **50**
- [Create a VMware Catalog Data Policy](#) on page **143**

# Configure Role-Based Access Control

[Role-Based Access Control Overview](#) on page **72**

[Configure Resource Pools](#) on page **74**

[Configure Roles](#) on page **77**

[Configure Accounts](#) on page **79**

[VMware Admin Role-Based Access Control Configuration](#) on page **81**

[NetApp ONTAP Admin Role-Based Access Control Configuration](#) on page **83**


[IBM Admin Role-Based Access Control Configuration](#) on page **85**


## 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.

**Note:** Users that register providers, such as storage devices, or add resources to IBM Spectrum Copy Data Management, such as policies 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 on the **Configure**  tab in the following order:

**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 74.

**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 Copy Data and Use Data policies, the role will determine how a user can interact with the policies. Permissions can be set to allow a user to create, view, and run the Copy Data and Use Data policies 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 77.

**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 79.



**RELATED TOPICS:**

- [Configure Resource Pools](#) on page **74**
- [Configure Roles](#) on page **77**
- [Configure Accounts](#) on page **79**
- [VMware Admin Role-Based Access Control Configuration](#) on page **81**
- [NetApp ONTAP Admin Role-Based Access Control Configuration](#) on page **83**
- [IBM Admin Role-Based Access Control Configuration](#) on page **85**

## 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 policies, so you must ensure a user running these policies has the correct permissions to interact with the hosts and policies 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 44.
- Add providers to IBM Spectrum Copy Data Management and associate them with a site. See [Register a Provider](#) on page 50.

### Add a Resource Pool




1. Click the **Configure**  tab. On the Views pane, select **Resource Pools** . The All Resource Pools pane opens.
2. In the All Resource Pools pane, click **New** . The New Resource Pool editor opens.
3. Click the **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 **2: Policies** tab. Select one or more policy types, individual custom policies, and schedules to include in the resource pool.
5. Click the **3: Reports** tab. Select one or more report types or individual reports to include in the resource pool.
6. Click the **4: Identities** tab. Select one or more keys and credentials to include in the resource pool.

7. Click the **5: Security** 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 Storage Workflows. 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.
8. Click the **6: 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.
9. Click the **7: 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 ECXALL resource pool cannot be edited.



1. Click the **Configure**  tab. On the Views pane, select **Resource Pools**  . The All Resource Pools pane opens.
2. In the All Resource Pools pane, 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 Association** to view the accounts that are associated with the resource pool. Re-assign your accounts to different resource pools before deleting.

**Note:** The ECXALL resource pool cannot be deleted.

1. Click the **Configure**  tab. On the Views pane, select **Resource Pools**  . The All Resource Pools pane opens.
2. In the All Resource Pools pane, 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 **77**.

**RELATED TOPICS:**

- [Configure Roles](#) on page **77**
- [Configure Accounts](#) on page **79**




## 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 Copy Data and Use Data policies, the role will determine how a user can interact with the policies. Permissions can be set to allow a user to create, view, and run the Copy Data and Use Data policies 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 44.
- Add providers to IBM Spectrum Copy Data Management and associate them with a site. See [Register a Provider](#) on page 50.
- 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 74.




### Add a Role

1. Click the **Configure**  tab. On the Views pane, select **Roles** . The All Roles pane opens.
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, policies, 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 Policy tabs as well as run and view reports. To set permissions, click **Click to select permissions**.
6. Continue adding features and setting associated permissions.
7. 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 **Roles** . The All Roles pane opens.
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 Association** to view the accounts that are associated with the role. Re-assign your accounts to different roles before deleting.

**Note:** The SYSADMIN and USER roles cannot be deleted.

1. Click the **Configure**  tab. On the Views pane, select **Roles** . The All Roles pane opens.
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 79.

### RELATED TOPICS:

- [Configure Resource Pools](#) on page 74
- [Configure Accounts](#) on page 79

## 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 44.
- Add providers to IBM Spectrum Copy Data Management and associate them with a site. See [Register a Provider](#) on page 50.
- 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 74.
- 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 77.

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

1. Click the **Configure**  tab. On the Views pane, select **Accounts** . The Accounts pane opens.
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 **Accounts** . The Accounts pane opens.
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 **Accounts**  . The Accounts pane opens.
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 **Accounts**  . The Accounts pane opens.
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 74 and [Configure Roles](#) on page 77.

#### RELATED TOPICS:

- [Configure Resource Pools](#) on page 74
- [Configure Roles](#) on page 77



## VMware Admin Role-Based Access Control Configuration

The VMware Admin Role manages VM resources, runs Copy and Use Data policies 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 Copy Data policies 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.

#### Policies Tab

3. Select the root level of all VMware related policies.
4. Select the root level of the IBM Spectrum Protect Snapshot Catalog policy.
5. Select the root Report policy.
6. Select the root All Storage Workflows policy.
7. Select the root All Schedules policy.

#### Reports Tab

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

#### Security 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 74
- [Configure Roles](#) on page 77
- [Configure Accounts](#) on page 79

# NetApp ONTAP Admin Role-Based Access Control Configuration

The NetApp ONTAP Admin Role manages NetApp ONTAP resources, runs Copy and Use Data policies 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 Copy and Use Data policies 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.

### Policies Tab

3. Select the root level of all NetApp related policies.
4. Select the root Report policy.
5. Select the root All Storage Workflows policy.
6. Select the root All Schedules policy.

### 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.

### Security 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 74
- [Configure Roles](#) on page 77
- [Configure Accounts](#) on page 79

## IBM Admin Role-Based Access Control Configuration

The IBM Admin Role manages IBM resources, runs Copy and Use Data policies 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 Copy and Use Data policies 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.

#### Policies Tab

3. Select the root level of IBM Catalog, IBM Copy, and IBM Use policies.
4. Select the root Report policy.
5. Select the root All Storage Workflows policy.
6. Select the root All Schedules policy.

#### Reports Tab

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

#### Security 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 74
- [Configure Roles](#) on page 77
- [Configure Accounts](#) on page 79

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

## BEFORE YOU BEGIN:

- Review [Best Practices for Configuring Tenants](#) on page 89.
- 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 74.
- 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 77.




## Add a Tenant

1. Click the **Configure**  tab. On the Views pane, select **Tenants** . The All Tenants pane opens.
2. In the All Tenants pane, 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 policies, 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 **Tenants**  . The All Tenants pane opens.
2. In the All Tenants pane, 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 policies and resources must be cleaned up through the Maintenance policy. The Maintenance policy removes resources and associated objects created by IBM Spectrum Copy Data Management when a policy 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 **Tenants**  . The All Tenants pane opens.
2. In the All Tenants pane, 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 89
- [Configure Resource Pools](#) on page 74





# 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**  pane. On the **Tenants**  pane, 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 policies and reports, but prevents them from viewing existing IBM Spectrum Copy Data Management policies:

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

## To assign resources to a tenant and allow the tenant users to create and run policies, 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**  pane. On the **Tenants**  pane, 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 policies 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 policies, 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 ECX\_All 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 **87**
- [Configure Resource Pools](#) on page **74**

# Configure Identities


[Identities Overview](#) on page **92**

[Add a Key](#) on page **93**

[Add a Credential](#) on page **96**

## 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**  tab.

### RELATED TOPICS:

- [Add a Key](#) on page **93**
- [Add a Credential](#) on page **96**

## 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 31.

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 *I will provide a keypair* method and register an associated provider




1. On an Oracle server, enter the `ssh-keygen` command and follow the instructions.
2. When asked to `Enter file in which to save the key`, enter a file and location such as the following: `/root/sshkey`.
3. In the `/root` location on the server entered in step 2, the file `sshkey.pub` contains the public key, which must be copied, pasted, and saved into the `authorized_keys` file after executing `cd ~/.ssh` while logged in as the Oracle user assigned to IBM Spectrum Copy Data Management.
4. In IBM Spectrum Copy Data Management, click the **Configure**  tab. On the Views pane, select **Identities** , then the **Keys** tab.
5. Click **New** . The Create Key dialog displays.
6. Select **SSH** as the key type and enter a key name in the Name field..
7. Select **I will provide a keypair** as the creation method.
8. In the `/root` location on the Oracle server entered in Step 2, the file `sshkey` contains the private key, which must be copied, pasted, and saved in the **Private Key** field in the Create Key dialog. Enter an optional comment.
9. Click **OK** to create the key. See the following steps to use this key to register an Oracle provider.
10. In IBM Spectrum Copy Data Management, click the **Configure**  tab. On the Views pane, select **Providers** . The Provider Browser opens.





11. Right-click **Oracle**  in the Provider Browser, then click **Register** . The Register Oracle Server dialog opens.
12. Select a Site, enter a Name and Host Address.
13. Select **Key** as the Authentication type. Enter the Oracle username, then select the key created in Step 5 in the Key field. Click **OK**.

### 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 **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 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 **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 **96**
- [Oracle Requirements](#) on page **31**
- [Register a Provider](#) on page **50**




## 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 **92**
- [Register a Provider](#) on page **50**



# Plan

[Plan Overview](#) on page **99**

[Configure Storage Workflows](#) on page **102**

## Schedules

- [Create a Schedule](#) on page **119**
- [Edit a Schedule](#) on page **121**
- [Delete a Schedule](#) on page **123**

## Catalog Data Policies

- [Create an Application Catalog Policy](#) on page **125**
- [Create a DellEMC Unity Catalog Data Policy](#) on page **127**
- [Create an IBM Spectrum Accelerate Catalog Data Policy](#) on page **129**
- [Create an IBM Spectrum Virtualize Catalog Data Policy](#) on page **131**
- [Create an IBM Spectrum Protect Snapshot Catalog Data Policy](#) on page **133**
- [Create a NetApp ONTAP Storage Catalog Data Policy](#) on page **135**
- [Create a NetApp ONTAP File Catalog Data Policy](#) on page **137**
- [Create a Pure Storage FlashArray Catalog Data Policy](#) on page **141**
- [Create a VMware Catalog Data Policy](#) on page **143**

## Copy Data Policies

- [Create an Application Copy Policy](#) on page **146**
- [Create a DellEMC Unity Copy Data Policy](#) on page **154**
- [Create an IBM Spectrum Accelerate Copy Data Policy](#) on page **157**
- [Create an IBM Spectrum Virtualize Copy Data Policy](#) on page **160**
- [Create a NetApp ONTAP Copy Data Policy](#) on page **164**
- [Create a Pure Storage FlashArray Copy Data Policy](#) on page **167**
- [Create a VMware Copy Data Policy](#) on page **170**
- [Create VMware Copy Data Policy Proxies](#) on page **175**

## Use Data Policies

- [Create an Application Use Policy](#) on page **179**
- [Create a DellEMC Unity Use Data Policy](#) on page **186**

- [Create an IBM Spectrum Accelerate Use Data Policy](#) on page **191**
- [Create an IBM Spectrum Virtualize Use Data Policy](#) on page **196**
- [Create a NetApp ONTAP Use Data Policy](#) on page **201**
- [Create a Pure Storage FlashArray Use Data Policy](#) on page **208**
- [Create a VMware Use Data Policy](#) on page **213**

[Using State and Status Arguments in Postscripts](#) on page **222**

[Create a Script Policy](#) on page **225**


[Create a Report Policy](#) on page **228**

[Edit a Policy](#) on page **230**

[Delete a Policy](#) on page **231**

[Maintenance Policy](#) on page **232**

## Plan Overview

Create policies, schedules, and jobs on the **Plan**  tab. You can also, edit, view and delete policies and schedules using this tab.

A policy is a user-defined set of tasks and rules. Once a policy is added to IBM Spectrum Copy Data Management, it can be combined with a schedule or trigger to create a job. There are several policy types including catalog, report, scripts, Copy Data and Use Data.

A schedule is a set of rules for triggering a policy. Create a schedule to apply to one or more jobs. Once a schedule is applied, the job sessions are run as defined by the parameters of the schedule.

A job is an execution of a policy. 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 policy so that you can easily add the schedule to the job definition in the New Policy editor.

### Policy Types

#### Catalog Data policies

Catalog policies interrogate storage systems to gather and record metadata about objects and files. There are two types of NetApp ONTAP catalog policies:

- NetApp ONTAP Storage Catalog Data policies, which are high-level objects.
- NetApp ONTAP File Catalog Data policies, which are low-level objects.

**Tip:** NetApp ONTAP File Catalog Data jobs might take a long time to run because they catalog at a file level. Consider creating NetApp ONTAP File Catalog Data policies that constrain the number of file system objects to catalog by limiting the number of storage systems and volume processed by a single policy. You can also stagger the scheduled run times of the policies to prevent them from running concurrently. Consider running NetApp ONTAP File Catalog Data policies less frequently than you run NetApp ONTAP Catalog Data policies.

For VMware, a catalog policy interrogates VMware hosts to gather and record information about vCenters, ESX servers, datastores, LUNs, and VMs, including VMDKs and datastores.

#### Copy Data and Use Data policies

IBM Spectrum Copy Data Management utilizes automated Copy Data Management workflows for replicating and intelligently reusing snapshots, vaults, and mirrors. Copy Data and Use Data policies offer control over testing and cloning use cases, instant recovery, and full disaster recovery. Through Copy Data and Use Data policies, you can:

- Copy data from a variety of storage providers to multiple locations.
- Reuse and recover resources from snapshots, vaults, and mirrors.
- Support use cases for automated data protection, recovery, DevOps, Dev/Test, data and database validation with data masking, through the use of automated Instant Access, Instant Virtualization, volume, and file restore functionalities.

### Report policies

A Report policy is a type of Analyze Data policy that summarizes information about cataloged providers and the data and other resources that reside on them.

### Script policies

A Script policy defines a set of commands to run on the IBM Spectrum Copy Data Management appliance. Use the script policy to add functionality to IBM Spectrum Copy Data Management. A script can consist of one or many commands, such as a shell script.

### Maintenance policy

The Maintenance policy removes resources and associated objects created by IBM Spectrum Copy Data Management when a policy 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 Storage Workflows](#) on page **102**
- [Create a Schedule](#) on page **119**
- [Create an Application Catalog Policy](#) on page **125**
- [Create a DellEMC Unity Catalog Data Policy](#) on page **127**
- [Create an IBM Spectrum Virtualize Catalog Data Policy](#) on page **131**
- [Create an IBM Spectrum Protect Snapshot Catalog Data Policy](#) on page **133**
- [Create a NetApp ONTAP Storage Catalog Data Policy](#) on page **135**
- [Create a NetApp ONTAP File Catalog Data Policy](#) on page **137**
- [Create a VMware Catalog Data Policy](#) on page **143**
- [Create an Application Copy Policy](#) on page **146**
- [Create a DellEMC Unity Copy Data Policy](#) on page **154**
- [Create an IBM Spectrum Virtualize Copy Data Policy](#) on page **160**
- [Create a NetApp ONTAP Copy Data Policy](#) on page **164**
- [Create a VMware Copy Data Policy](#) on page **170**
- [Create an Application Use Policy](#) on page **179**

- [Create a DellEMC Unity Use Data Policy](#) on page **186**
- [Create an IBM Spectrum Virtualize Use Data Policy](#) on page **196**
- [Create a NetApp ONTAP Use Data Policy](#) on page **201**
- [Create a VMware Use Data Policy](#) on page **213**
- [Create a Script Policy](#) on page **225**
- [Create a Report Policy](#) on page **228**
- [Monitor a Job Session](#) on page **237**

## Configure Storage Workflows

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

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

IBM Spectrum Copy Data Management Copy Data policies utilizing a VM Copy Storage Workflow can also be directed to Amazon Web Services (AWS) S3 destinations through an AWS Storage Gateway.

In a NetApp ONTAP Storage Workflow, 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:

- Register storage providers to be used in the Storage Workflows. Assign the storage providers to Sites. See [Register a Provider](#) on page 50 and [Add a Site](#) on page 44.
- If configuring an AWS Storage Workflow, 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 93.
- Create an account with the necessary permissions to create and run Copy Data policies. See [Role-Based Access Control Overview](#) on page 72.
- 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 50.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page 50.




### CONSIDERATIONS:

- Note that VMware Copy Data and Use Data policies only support vCenters or ESX hosts running vSphere 5.1 through 6.0.

### CONSIDERATIONS FOR VMWARE VIRTUAL VOLUMES:

- Storage Workflows that include virtual machines stored on virtual volume (VVOL) datastores through VM Copy sub-policies are supported. Replication is supported on the VM Copy 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 a Storage Workflow:


1. Click the **Plan**  tab. On the Views pane, select **Storage Workflows** . The All Storage Workflows pane opens.
2. In the All Storage Workflows pane, click **New** . The New Storage Workflow pane opens.
3. Select a type of workflow to create based on your storage provider. Select **AWS** to create an Amazon Web Services Copy Data workflow containing VM copies, **NetApp ONTAP** to create a NetApp ONTAP Copy Data workflow containing snapshots, VM copies, mirrors and vaults, **IBM** to create an IBM Copy Data workflow containing FlashCopies, Global Mirrors with Change Volumes, and VM Copies, or **DellEMC Unity** to create a DellEMC Unity Copy Data workflow containing snapshots and replication copies. VMware Copy Data policies support IBM, DellEMC Unity, and NetApp ONTAP Storage Workflows depending on your storage provider.
4. Enter a name and a meaningful description of the Storage Workflow.

### Configure AWS Storage Workflows

#### To add a VM Copy sub-policy to an AWS Storage Workflow:

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 Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field.

**Note:** Edits to the frequency and interval of a Storage Workflow apply to all associated policy schedules.

2. Click **Add VM Copy** .
3. In the VM Copy Destination pane select an AWS cloud destination from the list of available resources as the VM Copy 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 Copy 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 policy.

**Note:** Volume prefix labels must contain only alphanumeric characters and underscores. Labels cannot begin with numeric characters.

## Configure DellEMC Unity Storage Workflows

### To add a VM Copy sub-policy to a DellEMC Unity Storage Workflow:

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 Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field.

**Note:** Edits to the frequency and interval of a Storage Workflow apply to all associated policy schedules.

2. Click **Add VM Copy** .
3. In the VM Copy Destination pane select a DellEMC host destination from the list of available resources as the VM Copy 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 Copy 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 policy.

**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 policies utilizing the Storage Workflow.



**To add a snapshot sub-policy to a DellEMC Unity Storage Workflow:**

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 Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field.

**Note:** Edits to the frequency and interval of a Storage Workflow apply to all associated policy 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 policy.

**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 Storage Workflow:**

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 Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field.

**Note:** Edits to the frequency and interval of a Storage Workflow apply to all associated policy 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 policy.

**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 policy.

**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 policies utilizing the Storage Workflow.

**Configure IBM Spectrum Virtualize Storage Workflows****To add a VM Copy sub-policy to an IBM Spectrum Virtualize Storage Workflow:**

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 Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field.

**Note:** Edits to the frequency and interval of a Storage Workflow apply to all associated policy 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 Copy 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 Copy 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 policy.

**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 policy.

**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 policy. 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 Copy.

### Destination storage limit in GB / Destination volumes limit

Specify quotas for storage usage and the number of volume created on the destination for all policies utilizing the Storage Workflow.

### **To add a FlashCopy sub-policy to an IBM Spectrum Virtualize Storage Workflow:**


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 Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field.

**Note:** Edits to the frequency and interval of a Storage Workflow apply to all associated policy 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.

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 policy.

**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 Storage Workflow:*

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 Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field.

**Note:** Edits to the frequency and interval of a Storage Workflow apply to all associated policy 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 policy and cannot be edited after the policy 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 policy.

**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 policies utilizing the Storage Workflow.

## Configure IBM Spectrum Accelerate Storage Workflows

### *To add a VM Copy sub-policy to an IBM Spectrum Accelerate Storage Workflow:*

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 Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field.

**Note:** Edits to the frequency and interval of a Storage Workflow apply to all associated policy 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 Copy 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 Copy 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 policies 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.

#### **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 policy.

**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 policy.

**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 policy. 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 Copy.

**Destination storage limit in GB / Destination volumes limit**

Specify quotas for storage usage and the number of volume created on the destination for all policies utilizing the Storage Workflow.

**To add a snapshot sub-policy to an IBM Spectrum Accelerate Storage Workflow:**

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 Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field.

**Note:** Edits to the frequency and interval of a Storage Workflow apply to all associated policy 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 policies 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.

**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 policy.

**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 NetApp ONTAP Storage Workflows

### *To add a snapshot sub-policy to a NetApp ONTAP Storage Workflow:*

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 Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field.

**Note:** Edits to the frequency and interval of a Storage Workflow apply to all associated policy 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 policies 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 policy.

**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 Copy sub-policy to a NetApp ONTAP Storage Workflow:*

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 Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field.

**Note:** Edits to the frequency and interval of a Storage Workflow apply to all associated policy schedules.

2. Click **Add VM Copy** .

3. In the VM Copy Destination pane select an SVM from the list of available resources as the VM Copy 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 Copy 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 policies 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 policy.

**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.

**To add a mirror sub-policy to a NetApp ONTAP Storage Workflow:**

1. Select a snapshot, VM Copy, 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 policy and cannot be edited after the policy 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 policies utilizing the Storage Workflow.

### To add a vault sub-policy to a NetApp ONTAP Storage Workflow:

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 policy and cannot be edited after the policy 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 policies utilizing the Storage Workflow.

## Configure Pure Storage FlashArray Storage Workflows

### *To add a VM Copy sub-policy to a Pure Storage FlashArray Storage Workflow:*

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 Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field.

**Note:** Edits to the frequency and interval of a Storage Workflow apply to all associated policy schedules.

2. Click **Add VM Copy** .
3. In the VM Copy Destination pane select a Pure Storage FlashArray host destination from the list of available resources as the VM Copy destination.
4. In the Options pane set the VM Copy 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 policy.

**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 policy.

**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 policy. Available protocols include iSCSI and Fibre Channel.

### Full Copy Method

Select the full copy method. Available methods include Clone or VADP-based VM Copy.

### *To add a snapshot sub-policy to a Pure Storage FlashArray Storage Workflow:*

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 Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field.

**Note:** Edits to the frequency and interval of a Storage Workflow apply to all associated policy 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 policy.

**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 a Pure Storage FlashArray Storage Workflow:*

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 Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field.

**Note:** Edits to the frequency and interval of a Storage Workflow apply to all associated policy 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 policy.

**Note:** Snapshot labels must contain only alphanumeric characters and underscores.

5. When you are satisfied that the Storage Workflow-specific information is correct, click **Finish**. The Storage Workflow appears on the All Storage Workflows pane and can be applied to new and existing Copy Data policies.

### NEXT STEPS:

- Create a Copy Data policy that utilizes a Storage Workflow.

### RELATED TOPICS:

- [Create a DellEMC Unity Copy Data Policy](#) on page **154**
- [Create an IBM Spectrum Virtualize Copy Data Policy](#) on page **160**
- [Create a NetApp ONTAP Copy Data Policy](#) on page **164**
- [Create a Pure Storage FlashArray Copy Data Policy](#) on page **167**
- [Create a VMware Copy Data Policy](#) on page **170**

## Schedules

[Create a Schedule](#) on page **119**

[Edit a Schedule](#) on page **121**

[Delete a Schedule](#) on page **123**

## Create a Schedule





A schedule is a set of rules for triggering a policy. 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 121.

### To create a schedule:

1. Click the **Plan**  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 policy. See [Plan Overview](#) on page **99** and [Edit a Policy](#) on page **230**.

**RELATED TOPICS:**

- [Edit a Schedule](#) on page **121**
- [Delete a Schedule](#) on page **123**



## 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 119.

### To edit a schedule:

1. Click the **Plan**  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 policy. See [Plan Overview](#) on page **99** and [Edit a Policy](#) on page **230**.




**RELATED TOPICS:**

- [Create a Schedule](#) on page **119**
- [Delete a Schedule](#) on page **123**

## Delete a Schedule

Delete a schedule from the application if it is not used to trigger policies.

### To delete a schedule:

1. Click the **Plan**  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 **119**
- [Edit a Schedule](#) on page **121**

## Catalog Data Policies

[Create an Application Catalog Policy](#) on page **125**

[Create a DellEMC Unity Catalog Data Policy](#) on page **127**

[Create an IBM Spectrum Accelerate Catalog Data Policy](#) on page **129**

[Create an IBM Spectrum Virtualize Catalog Data Policy](#) on page **131**

[Create an IBM Spectrum Protect Snapshot Catalog Data Policy](#) on page **133**

[Create a NetApp ONTAP Storage Catalog Data Policy](#) on page **135**

[Create a NetApp ONTAP File Catalog Data Policy](#) on page **137**

[Create a Pure Storage FlashArray Catalog Data Policy](#) on page **141**

[Create a VMware Catalog Data Policy](#) on page **143**

## Create an Application Catalog Policy

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

An Application Catalog policy provides the framework to catalog and collect information about your Microsoft SQL and Oracle application servers.

Run an Application Catalog policy 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 a Microsoft SQL Server provider, you must also register associated vCenter providers.








### BEFORE YOU BEGIN:

- At least one application provider must be associated with an Application Catalog policy, such as an Oracle resource or VMware provider containing Microsoft application data. Before defining a catalog policy, add application providers. See [Register a Provider](#) on page 50.
- Credentials are required for cataloging application servers. See [Identities Overview](#) on page 92.
- Review Oracle requirements. See [Oracle Requirements](#) on page 31.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page 50.

### CONSIDERATIONS:

- One or more schedules might also be associated with a policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

### To create an Application Catalog policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **Application**  in the Catalog Data column. The Application Catalog Policy editor opens.
3. Select a **Microsoft on VMware** , **Microsoft SQL** , or **Oracle**  catalog data template.
4. Click the **1: Source** tab. From the list of available providers, select one or more providers containing application data that you wish to catalog. If creating an Application Catalog policy using a Microsoft template, registered VMware providers display in the Source pane. If creating an Application Catalog policy using an Oracle template, registered Oracle providers display in the source pane.

5. Click the **2: Notification** tab. Select the notification options for your policy.


### 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.

6. Click the **3: Schedule** tab. Select **Start job now** to create a policy 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

7. Click the **4: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**.

**Note:** If you selected the **Start job now** option, the job runs.

8. Click the **All Policies** tab. Your new policy appears in the policy list.

### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page 235.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page 237.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

### RELATED TOPICS:

- [Plan Overview](#) on page 99
- [Edit a Policy](#) on page 230
- [Delete a Policy](#) on page 231
- [Create a Schedule](#) on page 119

## Create a DellEMC Unity Catalog Data Policy

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

A DellEMC Unity Catalog Data policy 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 policy for cataloging.

Run a DellEMC Unity Catalog Data policy 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 Catalog Data policy. Before defining a catalog policy, add DellEMC Unity storage providers. See [Register a Provider](#) on page 50.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page 50.

### CONSIDERATIONS:

- One or more schedules might also be associated with a policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

### To create a DellEMC Unity Catalog Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **DellEMC Unity**  in the Catalog Data column. The DellEMC Unity Catalog Data Policy editor opens.
3. Click the **1: Source** tab. From the list of available providers, select one or more providers to catalog.
4. Click the **2: Options** tab. Select the options for your policy.

#### Number of catalog instances to keep

After a certain number of job runs for a given policy, older DellEMC Unity objects for that job are purged from the Catalog. Enter the number of job runs for which high-level DellEMC Unity objects are to be retained.

5. Click the **3: Notification** tab. Select the notification options for your policy.


#### 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.

- Click the **4: Schedule** tab. Select **Start job now** to create a policy 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

- Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**.

**Note:** If you selected the **Start job now** option, the job runs.

- Click the **All Policies** tab. Your new policy appears in the policy list.

### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page 235.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page 237.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

### RELATED TOPICS:

- [Plan Overview](#) on page 99
- [Edit a Policy](#) on page 230
- [Delete a Policy](#) on page 231
- [Create a Schedule](#) on page 119



## Create an IBM Spectrum Accelerate Catalog Data Policy

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

An IBM Spectrum Accelerate Catalog Data policy 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 policy for cataloging.

Run an IBM Spectrum Accelerate Catalog Data policy 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 Catalog Data policy. Before defining a catalog policy, add IBM Spectrum Accelerate storage providers. See [Register a Provider](#) on page 50.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page 50.

### CONSIDERATIONS:

- One or more schedules might also be associated with a policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

### To create an IBM Spectrum Accelerate Catalog Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **IBM Spectrum Accelerate**  in the Catalog Data column. The IBM Spectrum Accelerate Catalog Data Policy editor opens.
3. Click the **1: Source** tab. From the list of available providers, select one or more providers to catalog.
4. Click the **2: Options** tab. Select the options for your policy.

#### Number of catalog instances to keep

After a certain number of job runs for a given policy, older IBM Spectrum Accelerate objects for that job are purged from the Catalog. Enter the number of job runs for which high-level IBM Spectrum Accelerate objects are to be retained.

5. Click the **3: Notification** tab. Select the notification options for your policy.


#### 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.

- Click the **4: Schedule** tab. Select **Start job now** to create a policy 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

- Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**.

**Note:** If you selected the **Start job now** option, the job runs.

- Click the **All Policies** tab. Your new policy appears in the policy list.

### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page 235.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page 237.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

### RELATED TOPICS:

- [Plan Overview](#) on page 99
- [Edit a Policy](#) on page 230
- [Delete a Policy](#) on page 231
- [Create a Schedule](#) on page 119

## Create an IBM Spectrum Virtualize Catalog Data Policy

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

An IBM Spectrum Virtualize Catalog Data policy 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 policy for cataloging.

Run an IBM Spectrum Virtualize Catalog Data policy 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 Catalog Data policy. Before defining a catalog policy, add IBM storage providers. See [Register a Provider](#) on page 50.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page 50.

### CONSIDERATIONS:

- IBM providers utilize port 22 for communication with IBM Spectrum Copy Data Management.
- One or more schedules might also be associated with a policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

### To create an IBM Spectrum Virtualize Catalog Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **IBM Spectrum Virtualize**  in the Catalog Data column. The IBM Spectrum Virtualize Catalog Data Policy editor opens.
3. Click the **1: Source** tab. From the list of available providers, select one or more providers to catalog.
4. Click the **2: Options** tab. Select the options for your policy.

#### Number of catalog instances to keep

After a certain number of job runs for a given policy, older IBM objects for that job are purged from the Catalog. Enter the number of job runs for which high-level IBM objects are to be retained.

5. Click the **3: Notification** tab. Select the notification options for your policy.


#### 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.

6. Click the **4: Schedule** tab. Select **Start job now** to create a policy 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

7. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**.

**Note:** If you selected the **Start job now** option, the job runs.

8. Click the **All Policies** tab. Your new policy appears in the policy list.

### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page 235.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page 237.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

### RELATED TOPICS:

- [Plan Overview](#) on page 99
- [Edit a Policy](#) on page 230
- [Delete a Policy](#) on page 231
- [Create a Schedule](#) on page 119

## Create an IBM Spectrum Protect Snapshot Catalog Data Policy

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

An IBM Spectrum Protect Snapshot Catalog Data policy provides the framework to catalog and collect information about high-level IBM objects.

Run an IBM Spectrum Protect Snapshot Catalog Data policy 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 Catalog Data policy. Before defining a catalog policy, add an IBM Spectrum Protect Snapshot provider. See [Register a Provider](#) on page 50.
- 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 policy, add SMTP resources. See [Register a Provider](#) on page 50.

### CONSIDERATIONS:

- One or more schedules might also be associated with a policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

### To create an IBM Spectrum Protect Snapshot Catalog Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **IBM Spectrum Protect Snapshot**  in the Catalog Data column. The IBM Spectrum Protect Snapshot Catalog Data Policy editor opens.
3. Click the **1: Source** tab. From the list of available providers, select one or more providers to catalog.
4. Click the **2: Notification** tab. Select the notification options for your policy.


#### 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.

5. Click the **3: Schedule** tab. Select **Start job now** to create a policy 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

6. Click the **4: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**.

**Note:** If you selected the **Start job now** option, the job runs.

7. Click the **All Policies** tab. Your new policy appears in the policy list.

### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page 235.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page 237.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

### RELATED TOPICS:

- [Plan Overview](#) on page 99
- [Edit a Policy](#) on page 230
- [Delete a Policy](#) on page 231
- [Create a Schedule](#) on page 119

## Create a NetApp ONTAP Storage Catalog Data Policy

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

A NetApp ONTAP Storage Catalog Data policy 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 policy for cataloging. NetApp ONTAP 7-Mode and Cluster-Mode are both supported.

Run a NetApp ONTAP Storage Catalog Data policy 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 Catalog Data policy.





### BEFORE YOU BEGIN:

- At least one NetApp ONTAP provider must be associated with a NetApp ONTAP File Catalog Data policy. Before defining a catalog policy, add NetApp ONTAP providers. See [Register a Provider](#) on page 50.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page 50.

### CONSIDERATIONS:

- One or more schedules might also be associated with a policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

### To create a NetApp ONTAP Storage Catalog Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **NetApp ONTAP Storage**  in the Catalog Data column. The NetApp ONTAP Storage Catalog Data Policy editor opens.
3. Click the **1: Source** tab. From the list of available providers, select one or more providers to catalog.
4. Click the **2: Options** tab. Select the options for your policy.

#### 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 policy, older NetApp ONTAP objects for that job are purged from the Catalog. Enter the number of job runs for which high-level NetApp ONTAP objects are to be retained.

5. Click the **3: Notification** tab. Select the notification options for your policy.


#### 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.

6. Click the **4: Schedule** tab. Select **Start job now** to create a policy 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

7. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**.

**Note:** If you selected the **Start job now** option, the job runs.

8. Click the **All Policies** tab. Your new policy appears in the policy list.

#### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page 235.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page 237.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

#### RELATED TOPICS:

- [Plan Overview](#) on page 99
- [Create a NetApp ONTAP File Catalog Data Policy](#) on page 137
- [Edit a Policy](#) on page 230
- [Delete a Policy](#) on page 231
- [Create a Schedule](#) on page 119



## Create a NetApp ONTAP File Catalog Data Policy

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

A NetApp ONTAP File Catalog Data policy 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 policy 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 Catalog Data policy 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 Catalog Data policy.

**Tip:** NetApp ONTAP File Catalog Data jobs might take a long time to run because they catalog at a file level. Consider creating NetApp ONTAP File Catalog Data policies that constrain the number of file system objects to catalog by limiting the number of storage systems and volume processed by a single policy. You can also stagger the scheduled run times of the policies to prevent them from running concurrently. Consider running NetApp ONTAP File Catalog Data policies less frequently than you run NetApp ONTAP Catalog Data policies.

### BEFORE YOU BEGIN:





- At least one NetApp ONTAP provider must be associated with a NetApp ONTAP File Catalog Data policy. Before defining a catalog policy, add NetApp ONTAP providers. See [Register a Provider](#) on page 50.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page 50.

### CONSIDERATIONS:

- One or more schedules might also be associated with a policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

**Best Practice:** Create a schedule before creating a policy so that you can easily add the schedule to the job definition in the New Policy editor.

**To create a NetApp ONTAP File Catalog Data policy:**

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **NetApp ONTAP Files**  in the Catalog Data column. The NetApp ONTAP File Catalog Data Policy editor opens.

3. Click the **1: Source** tab. From the list of available providers, 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.
4. Click the **2: Options** tab. Select the options for your policy.

**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 policy, older low-level NetApp ONTAP objects for that job are purged from the Catalog. 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 policy is purged from the IBM Spectrum Copy Data Management Catalog 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 Catalog before newer data is cataloged through a NetApp ONTAP File Catalog Data policy.

**Condense catalog after failed run**

Select this option to purge NetApp ONTAP objects from the IBM Spectrum Copy Data Management Catalog for the unsuccessful run of the NetApp ONTAP File Catalog Data policy.

**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 Catalog. See the Filewalker documentation for more information.

5. Click the **3: Notification** tab. Select the notification options for your policy.


### 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.

6. Click the **4: Schedule** tab. Select **Start job now** to create a policy 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

7. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**.

**Note:** If you selected the **Start job now** option, the job runs.

8. Click the **All Policies** tab. Your new policy appears in the policy list.

### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page **235**.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

### RELATED TOPICS:

- [Plan Overview](#) on page **99**
- [Create a NetApp ONTAP Storage Catalog Data Policy](#) on page **135**
- [Edit a Policy](#) on page **230**
- [Delete a Policy](#) on page **231**

- [Create a Schedule](#) on page **119**
- [Search and Filter Guidelines](#) on page **404**

## Create a Pure Storage FlashArray Catalog Data Policy

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

A Pure Storage FlashArray Catalog Data policy 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 policy for cataloging.

Run a Pure Storage FlashArray Catalog Data policy 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 Catalog Data policy. Before defining a catalog policy, add Pure Storage providers. See [Register a Provider](#) on page 50.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page 50.

### CONSIDERATIONS:

- One or more schedules might also be associated with a policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

### To create a Pure Storage FlashArray Catalog Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **Pure Storage FlashArray**  in the Catalog Data column. The Pure Storage FlashArray Catalog Data Policy editor opens.
3. Click the **1: Source** tab. From the list of available providers, select one or more providers to catalog.
4. Click the **2: Options** tab. Select the options for your policy.

#### Number of catalog instances to keep

After a certain number of job runs for a given policy, older Pure Storage objects for that job are purged from the Catalog. Enter the number of job runs for which high-level Pure Storage objects are to be retained.

5. Click the **3: Notification** tab. Select the notification options for your policy.


#### 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.

6. Click the **4: Schedule** tab. Select **Start job now** to create a policy 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

7. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**.

**Note:** If you selected the **Start job now** option, the job runs.

8. Click the **All Policies** tab. Your new policy appears in the policy list.

### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page 235.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page 237.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

### RELATED TOPICS:

- [Plan Overview](#) on page 99
- [Edit a Policy](#) on page 230
- [Delete a Policy](#) on page 231
- [Create a Schedule](#) on page 119

## Create a VMware Catalog Data Policy

Cataloging objects located on a provider is required for browsing, searching, and reporting in IBM Spectrum Copy Data Management.

A VMware Catalog Data policy provides the framework to catalog and collect information about VMware objects. You can select one or more VMware providers in a single policy for cataloging.

Run a VMware Catalog Data policy 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 Catalog Data policy. Before defining a catalog policy, add VMware providers. See [Register a Provider](#) on page 50.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page 50.

### CONSIDERATIONS:

- One or more schedules might also be associated with a policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

### To create a VMware Catalog Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **VMware**  in the Catalog Data column. The VMware Catalog Data Policy editor opens.
3. Click the **1: Source** tab. From the list of available providers, select one or more providers to catalog.
4. Click the **2: Options** tab. Select the options for your policy.

#### 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 policy, older VMware objects for that job are purged from the Catalog. Enter the number of job runs for which VMware objects are to be retained.

5. Click the **3: Notification** tab. Select the notification options for your policy.


#### 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.

6. Click the **4: Schedule** tab. Select **Start job now** to create a policy 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

7. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**.

**Note:** If you selected the **Start job now** option, the job runs.

8. Click the **All Policies** tab. Your new policy appears in the policy list.

### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page **235**.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

### RELATED TOPICS:

- [Plan Overview](#) on page **99**
- [Edit a Policy](#) on page **230**
- [Delete a Policy](#) on page **231**
- [Create a Schedule](#) on page **119**



## Copy Data Policies

[Create an Application Copy Policy](#) on page **146**

[Create a DellEMC Unity Copy Data Policy](#) on page **154**

[Create an IBM Spectrum Accelerate Copy Data Policy](#) on page **157**

[Create an IBM Spectrum Virtualize Copy Data Policy](#) on page **160**

[Create a NetApp ONTAP Copy Data Policy](#) on page **164**

[Create a Pure Storage FlashArray Copy Data Policy](#) on page **167**

[Create a VMware Copy Data Policy](#) on page **170**

[Create VMware Copy Data Policy Proxies](#) on page **175**

## Create an Application Copy Policy

IBM Spectrum Copy Data Management provides application database copy management through application-consistent copy 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 and 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 copies only of eligible databases. To be eligible for copy, application databases must reside on supported storage platforms.

The following options are available in the Application Copy policy:

**RMAN Integration** - Oracle-based policies. 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 copies in the RMAN recovery catalog, enabling database administrators to leverage RMAN for verification and advanced recovery.

**Data Masking** - Oracle-based policies. 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** - Oracle and Microsoft SQL-based policies. The log backup feature enables continuous copy of Archive logs to a specified destination. In Oracle-based Application Copy policies, 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 an Application Catalog policy that includes the providers you wish to copy. See [Create an Application Catalog Policy](#) on page 125.
- Configure a Storage Workflow. See [Configure Storage Workflows](#) on page 102.
- Review Oracle requirements. See [Oracle Requirements](#) on page 31.

### 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.
- Note that an automated disk mount of an IBM Spectrum Copy Data Management copy on LVM storage must not have the original disk for that copy still present and mounted while the lvm2-lvmetad service is active on the system. If Oracle data resides on LVM volumes,

you must stop and disable the lvm2-lvmetad service before running Application Copy or Use jobs. To disable the lvm2-lvmetad, 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 Application Copy Policy 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.
- 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 copy. 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 policies 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.







#### **MICROSOFT SQL SERVER CONSIDERATIONS:**

- IBM Spectrum Copy Data Management supports one Microsoft SQL database per SQL-based Application Copy policy, therefore you must avoid protecting a SQL database through multiple Application Copy policies.
- In the Log Backup step of a SQL Application Copy policies, there is no mechanism to ensure retention or log cleanup, which can result in a destination mount point exceeding its capacity.
- Note that IBM Spectrum Copy Data Management does not support log backup of Simple recovery models.

#### **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 copy job ran.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page 50.
- One or more schedules might also be associated with a policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

### **To create an Oracle Database Application Copy policy:**

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, click **New** , then select **Application**  in the Copy Data column. The Application Copy Policy editor opens.
3. Select an **Oracle**  template.
4. Click the **1: Source** tab. From the list of available sites select one or more providers to copy as defined in the associated Storage Workflows. 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 is running in NOARCHIVELOG mode, or the database files, control files, or redo log files are stored on unsupported storage.
5. Click the **2. Storage Workflow** tab, then select a Storage Workflow that meets your copy data criteria.
6. Set a time to run the Storage Workflow with the **Specify Activation Time** trigger, then click **Add** to add it to the policy. To run the Storage Workflow on demand through the **Monitor**  tab, select **Manual**.

If configuring more than one Storage Workflow in a policy, use the **Same as workflow** option to trigger multiple Storage Workflows to run at the same time.

7. Repeat the steps above to add additional Storage Workflows to the policy.
8. Click the **3. Options** tab. Set the policy options.

#### **Maximum Concurrent Tasks**

Set the maximum amount of concurrent transfers between the source and the destination.

#### **Record copies in RMAN local repository**

Enable to create a local copy of the Recovery Manager (RMAN) catalog during the running of the Application Copy policy. 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 Application Copy Policy utilizing the **Record copies in RMAN recovery catalog** feature.

### Enable Policy-Level Scripts

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the `/data/userscripts` directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: `/data/userscripts/<name of script>`. For information about script return codes, see [Return Code Reference](#) on page 409.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

### Enable Policy-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot subpolicy runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. Enter individual script commands or the location of the prescript or postscript.

A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the `/data/userscripts` directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: `/data/userscripts/<name of script>`.

**\_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 policy. 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 policy if a command in any of the scripts associated with the policy fails.

9. Click the **4. Log Backup** tab. 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.

If multiple databases are selected for Copy Data, 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 policy, and a single destination directory named /logbackup is defined in the policy, then you must create separate disks for both servers and mount them both to /logbackup on the individual servers.

If the **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 Application Copy policy.

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 policy. After a successful copy, logs older than that copy 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 policy, 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.

10. Click the **5. Data Masking** tab. If enabled, IBM Spectrum Copy Data Management mounts snapshot copies of the protected database onto a user-specified staging server. Select resources to be masked from the list of available databases, select a copy 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.

IBM Spectrum Copy Data Management spins up a clone of the database on the staging server, 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 Application Use policies.

**Note:** Oracle homes selected to be protected must be different from the Oracle home where masking takes place in the policy.


11. Click the **6: Notification** tab. Select the policy 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.

- Click the **7: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.

### To create a Microsoft SQL Server Application Copy policy:

- Click the **Plan**  tab. On the Views pane, select **Policies** .
- Click the **All Policies** tab, click **New** , then select **Application**  in the Copy Data column. The Application Copy Policy editor opens.
- Select **Microsoft SQL** , then select a **Standalone** or **Always On** workflow template.
- Click the **1: Source** tab. From the list of available sites select a provider to copy as defined in the associated Storage Workflows. 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 is running in NOARCHIVELOG mode, or the database files, control files, or redo log files are stored on unsupported storage.
- Click the **2. Storage Workflow** tab, then select a Storage Workflow that meets your copy data criteria.
- Set a time to run the Storage Workflow with the **Specify Activation Time** trigger, then click **Add** to add it to the policy. To run the Storage Workflow on demand through the **Monitor**  tab, select **Manual**.

If configuring more than one Storage Workflow in a policy, use the **Same as workflow** option to trigger multiple Storage Workflows to run at the same time.

- Repeat the steps above to add additional Storage Workflows to the policy.
- Click the **3. Options** tab. Set the policy options.

### Maximum Concurrent Tasks

Set the maximum amount of concurrent transfers between the source and the destination.

### Maximum Concurrent Snapshots on ESX

Set the maximum number of concurrent snapshots on the vCenter.

### Enable Policy-Level Scripts

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the `/data/userscripts` directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: `/data/userscripts/<name of script>`. For information about script return codes, see [Return Code Reference](#) on page 409.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

### Enable Policy-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot subpolicy runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. Enter individual script commands or the location of the prescript or postscript.

A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the `/data/userscripts` directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: `/data/userscripts/<name of script>`.

**\_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 policy. 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 policy if a command in any of the scripts associated with the policy fails.

9. Click the **4. Log Backup** tab. Select **Backup logs** to back up database logs and protect 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 policies 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 Copy Data, then each of the servers hosting the databases must have their Destination Mount Points set individually. For example, if two databases from Server A and Server B are added to the same policy, and a single mount point named `/logbackup` is defined in the policy, then you must create separate disks for both servers 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.

10. Click the **5: Notification** tab. Select the policy 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. Click the **6: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.

**NEXT STEPS:**

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page **235**.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Create an Application Use policy. See [Create an Application Use Policy](#) on page **179**.

**RELATED TOPICS:**

- [Plan Overview](#) on page **99**
- [Edit a Policy](#) on page **230**
- [Delete a Policy](#) on page **231**
- [Create a Schedule](#) on page **119**
- [Create an Application Use Policy](#) on page **179**

## Create a DellEMC Unity Copy Data Policy

Copy DellEMC Unity data with snapshots, replication, and VM Copies using a DellEMC Unity Copy Data policy. The RPO and copy data parameters are defined in a Storage Workflow, which is then applied to the Copy Data policy along with a specified activation time to meet your copy data criteria. Supported sources include DellEMC Unity storage systems.






### BEFORE YOU BEGIN:

- Create and run a DellEMC Unity Catalog Data policy that includes the providers you wish to copy. See [Create a DellEMC Unity Catalog Data Policy](#) on page 127.
- Configure a DellEMC Unity Storage Workflow. See [Configure Storage Workflows](#) on page 102.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page 50.

### 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 policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

### To create a DellEMC Unity Copy Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, click **New** , then select **DellEMC Unity**  in the Copy Data column. The DellEMC Unity Copy Data Policy editor opens.
3. Click the **1: Source** tab. From the drop-down menu select **LUNs** or **File Systems**. From the list of available sites, select one or more providers to copy as defined in the associated Storage Workflows.
4. Click the **2. Storage Workflow** tab, then select a Storage Workflow that meets your copy data criteria.
5. Set a time to run the Storage Workflow with the **Specify activation time** trigger, then click **Add** to add it to the policy. To run the Storage Workflow on demand through the **Monitor**  tab, select **Manual**.

If configuring more than one Storage Workflow in a policy, select the **Same as workflow** option to trigger multiple Storage Workflows to run concurrently.

**Note:** Only Storage Workflows with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating a Storage Workflow.

6. Repeat the steps above to add additional Storage Workflows to the policy.

7. Click the **3. Options** tab. Set the policy options.

### Maximum Concurrent Tasks

Set the maximum amount of concurrent transfers between the source and the destination.

### Enable Policy-Level Scripts

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>. For information about script return codes, see [Return Code Reference](#) on page 409.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.


8. Click the **4: Notification** tab. Select the policy 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. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab

### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page 235.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page 237.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Use the Catalog Browse feature to review the recovery point. See [Browse Catalog](#) on page 252.
- Create a DellEMC Unity Use Data policy. See [Create a DellEMC Unity Use Data Policy](#) on page 186.

**RELATED TOPICS:**

- [Plan Overview](#) on page **99**
- [Edit a Policy](#) on page **230**
- [Delete a Policy](#) on page **231**
- [Create a Schedule](#) on page **119**
- [Create a DellEMC Unity Use Data Policy](#) on page **186**

## Create an IBM Spectrum Accelerate Copy Data Policy

Copy IBM Spectrum Accelerate data with FlashCopies, Global Mirrors, and VM Copies using an IBM Spectrum Accelerate Copy Data policy. The RPO and copy data parameters are defined in a Storage Workflow, which is then applied to the Copy Data policy along with a specified activation time to meet your copy data criteria.






### BEFORE YOU BEGIN:

- Create and run an IBM Spectrum Accelerate Catalog Data policy that includes the providers you wish to copy. See [Create an IBM Spectrum Accelerate Catalog Data Policy](#) on page **129**.
- Configure an IBM Spectrum Accelerate Storage Workflow. See [Configure Storage Workflows](#) on page **102**.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page **50**.

### 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 policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page **119**.

### To create an IBM Spectrum Accelerate Copy Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, click **New** , then select **IBM Spectrum Accelerate**  in the Copy Data column. The IBM Spectrum Accelerate Copy Data Policy editor opens.
3. Click the **1: Source** tab. From the list of available sites select one or more providers to copy as defined in the associated Storage Workflows.
4. Click the **2. Storage Workflow** tab, then select a Storage Workflow that meets your copy data criteria.
5. Set a time to run the Storage Workflow with the **Specify activation time** trigger, then click **Add** to add it to the policy. To run the Storage Workflow on demand through the **Monitor**  tab, select **Manual**.

If configuring more than one Storage Workflow in a policy, select the **Same as workflow** option to trigger multiple Storage Workflows to run concurrently.

**Note:** Only Storage Workflows with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating a Storage Workflow.

- Repeat the steps above to add additional Storage Workflows to the policy.
- Click the **3. Options** tab. Set the policy 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 Storage Workflow 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 policy, a Consistency Group will be created for each provider. Consistency Groups are named based on the prefix provided during policy creation plus the policy name.

### Enable Policy-Level Scripts

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>. For information about script return codes, see [Return Code Reference](#) on page 409.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

### Enable Policy-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot subpolicy runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. Enter individual script commands or the location of the prescript or postscript.

A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>.

**\_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 policy. 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 policy if a command in any of the scripts associated with the policy fails.


8. Click the **4: Notification** tab. Select the policy 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. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab

#### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page **235**.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Use the Catalog Browse feature to review the recovery point. See [Browse Catalog](#) on page **252**.
- Create an IBM Spectrum Accelerate Use Data policy. See [Create an IBM Spectrum Accelerate Use Data Policy](#) on page **191**.

#### RELATED TOPICS:

- [Plan Overview](#) on page **99**
- [Edit a Policy](#) on page **230**
- [Delete a Policy](#) on page **231**
- [Create a Schedule](#) on page **119**
- [Create an IBM Spectrum Accelerate Use Data Policy](#) on page **191**

## Create an IBM Spectrum Virtualize Copy Data Policy

Copy IBM data with FlashCopies, Global Mirrors, and VM Copies using an IBM Spectrum Virtualize Copy Data policy. The RPO and copy data parameters are defined in a Storage Workflow, which is then applied to the Copy Data policy 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 Catalog Data policy that includes the providers you wish to copy. See [Create an IBM Spectrum Virtualize Catalog Data Policy](#) on page 131.
- Configure an IBM Storage Workflow. See [Configure Storage Workflows](#) on page 102.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page 50.

### 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 policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

### To create an IBM Spectrum Virtualize Copy Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, click **New** , then select **IBM Spectrum Virtualize**  in the Copy Data column. The IBM Spectrum Virtualize Copy Data Policy editor opens.
3. Click the **1: Source** tab. From the list of available sites select one or more providers to copy as defined in the associated Storage Workflows. To exclude Flash Copies from the list of sources, select **Exclude Flash Copies**.
4. Click the **2. Storage Workflow** tab, then select a Storage Workflow that meets your copy data criteria.
5. Set a time to run the Storage Workflow with the **Specify activation time** trigger, then click **Add** to add it to the policy. To run the Storage Workflow on demand through the **Monitor**  tab, select **Manual**.



If configuring more than one Storage Workflow in a policy, select the **Same as workflow** option to trigger multiple Storage Workflows to run concurrently.

**Note:** Only Storage Workflows with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating a Storage Workflow.

6. Repeat the steps above to add additional Storage Workflows to the policy.
7. Click the **3. Options** tab. Set the policy options.

### Skip the Flash Copy Target Volumes

Select this option to ensure FlashCopy target volumes are excluded from policies associated with the Storage Workflow.

### 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 Storage Workflow 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 policy, a Consistency Group will be created for each provider. Consistency Groups are named based on the prefix provided during policy creation plus the policy name.

### Enable Policy-Level Scripts

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the `/data/userscripts` directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: `/data/userscripts/<name of script>`. For information about script return codes, see [Return Code Reference](#) on page 409.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

### Enable Policy-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot subpolicy runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. Enter individual script commands or the location of the prescript or postscript.

A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the `/data/userscripts` directory on the IBM Spectrum Copy Data Management appliance. To

invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: `/data/userscripts/<name of script>`.

**\_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 policy. 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 policy if a command in any of the scripts associated with the policy fails.


8. Click the **4: Notification** tab. Select the policy 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. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab

#### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page **235**.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Use the Catalog Browse feature to review the recovery point. See [Browse Catalog](#) on page **252**.
- Create an IBM Spectrum Virtualize Use Data policy. See [Create an IBM Spectrum Virtualize Use Data Policy](#) on page **196**.

#### RELATED TOPICS:

- [Plan Overview](#) on page **99**
- [Edit a Policy](#) on page **230**
- [Delete a Policy](#) on page **231**

- [Create a Schedule](#) on page **119**
- [Create an IBM Spectrum Virtualize Use Data Policy](#) on page **196**

## Create a NetApp ONTAP Copy Data Policy

Copy NetApp ONTAP data with snapshots using a NetApp ONTAP Copy Data policy. A NetApp ONTAP Copy Data policy consists of snapshot, mirror, and vault sub-policies defined in a Storage Workflow, 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.


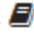



### BEFORE YOU BEGIN:

- Create and run a NetApp ONTAP Storage Catalog Data policy that includes the providers you wish to copy. See [Create a NetApp ONTAP Storage Catalog Data Policy](#) on page **135**.
- Configure a NetApp ONTAP Storage Workflow. See [Configure Storage Workflows](#) on page **102**.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page **50**.

### CONSIDERATIONS:

- Note that cloned volumes will not be replicated through a Copy Data policy.
- Note that snapshot postscript functionality applies only to NetApp ONTAP storage snapshot subpolicies.
- One or more schedules might also be associated with a policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page **119**.

### To create a NetApp ONTAP Copy Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, click **New** , then select **NetApp ONTAP**  in the Copy Data column. The NetApp ONTAP Copy Data Policy editor opens.
3. Click the **1: Source** tab. From the list of available sites select one or more providers to copy as defined in the associated Storage Workflows.
4. Click the **2. Storage Workflow** tab, then select a Storage Workflow that meets your copy data criteria.
5. Set a time to run the Storage Workflow with the **Specify Activation Time** trigger, then click **Add** to add it to the policy. To run the Storage Workflow on demand through the **Monitor**  tab, select **Manual**.
6. Repeat the steps above to add additional Storage Workflows to the policy.
7. Click the **3. Options** tab. Set the policy options.

### Maximum Concurrent Tasks

Set the maximum amount of concurrent transfers between the source and the destination.

### Enable Policy-Level Scripts

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the `/data/userscripts` directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: `/data/userscripts/<name of script>`. For information about script return codes, see [Return Code Reference](#) on page 409.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

### Enable Policy-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot subpolicy runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. Enter individual script commands or the location of the prescript or postscript.

A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the `/data/userscripts` directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: `/data/userscripts/<name of script>`.

**\_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 policy. 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 policy if a command in any of the scripts associated with the policy fails.


8. Click the **4: Notification** tab. Select the policy 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. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab

### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page **235**.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Use the Catalog Browse feature to review the recovery point. See [Browse Catalog](#) on page **252**.
- Create a NetApp ONTAP Use Data policy. See [Create a NetApp ONTAP Use Data Policy](#) on page **201**.

**RELATED TOPICS:**

- [Plan Overview](#) on page **99**
- [Edit a Policy](#) on page **230**
- [Delete a Policy](#) on page **231**
- [Create a Schedule](#) on page **119**
- [Create a NetApp ONTAP Use Data Policy](#) on page **201**

## Create a Pure Storage FlashArray Copy Data Policy

Copy Pure Storage data with snapshots and replications using a Pure Storage FlashArray Copy Data policy. The RPO and copy data parameters are defined in a Storage Workflow, which is then applied to the Copy Data policy along with a specified activation time to meet your copy data criteria.






### BEFORE YOU BEGIN:

- Create and run a Pure Storage FlashArray Catalog Data policy that includes the providers you wish to copy. See [Create a Pure Storage FlashArray Catalog Data Policy](#) on page **141**.
- Configure a Pure Storage FlashArray Storage Workflow. See [Configure Storage Workflows](#) on page **102**.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page **50**.

### CONSIDERATIONS:

- One or more schedules might also be associated with a policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page **119**.

### To create a Pure Storage FlashArray Copy Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, click **New** , then select **Pure Storage FlashArray**  in the Copy Data column. The Pure Storage FlashArray Copy Data Policy editor opens.
3. Click the **1: Source** tab. From the list of available sites select one or more providers to copy as defined in the associated Storage Workflows.
4. Click the **2. Storage Workflow** tab, then select a Storage Workflow that meets your copy data criteria.
5. Set a time to run the Storage Workflow with the **Specify activation time** trigger, then click **Add** to add it to the policy. To run the Storage Workflow on demand through the **Monitor**  tab, select **Manual**.

If configuring more than one Storage Workflow in a policy, select the **Same as workflow** option to trigger multiple Storage Workflows to run concurrently.

**Note:** Only Storage Workflows with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating a Storage Workflow.

6. Repeat the steps above to add additional Storage Workflows to the policy.
7. Click the **3. Options** tab. Set the policy options.

### 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 Copy Data functions on the entire group. If the associated Storage Workflow contains different subpolicy types, a separate Consistency Group will be created for each copy type. Note that if more than one Pure Storage FlashArray provider is selected in the policy, a Consistency Group will be created for each provider. Consistency Groups are named based on the prefix provided during policy creation plus the policy name.

### Enable Policy-Level Scripts

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>. For information about script return codes, see [Return Code Reference](#) on page 409.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

### Enable Policy-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot subpolicy runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. Enter individual script commands or the location of the prescript or postscript.

A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>.

**\_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 policy. 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 policy if a command in any of the scripts associated with the policy fails.

8. Click the **4: Notification** tab. Select the policy 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. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab

**NEXT STEPS:**

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page **235**.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Use the Catalog Browse feature to review the recovery point. See [Browse Catalog](#) on page **252**.
- Create a Pure Storage FlashArray Use Data policy. See [Create a Pure Storage FlashArray Use Data Policy](#) on page **208**.

**RELATED TOPICS:**

- [Plan Overview](#) on page **99**
- [Edit a Policy](#) on page **230**
- [Delete a Policy](#) on page **231**
- [Create a Schedule](#) on page **119**
- [Create a Pure Storage FlashArray Use Data Policy](#) on page **208**

## Create a VMware Copy Data Policy

Copy VMware data including virtual machines, datastores, folders, vApps, and datacenters with snapshots using a VMware Copy Data policy. A VMware Copy Data policy consists of snapshot, mirror, and vault sub-policies defined in a Storage Workflow, 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 Copy Data policies support IBM, DellEMC Unity, and NetApp ONTAP Storage Workflows.

### BEFORE YOU BEGIN:

- Create and run a VMware Catalog Data policy that includes the providers you wish to copy. See [Create a VMware Catalog Data Policy](#) on page **143**.
- Configure a Storage Workflow. See [Configure Storage Workflows](#) on page **102**.
- 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 policy, add SMTP resources. See [Register a Provider](#) on page **50**.

### CONSIDERATIONS:






- Note that VMware Copy Data and Use Data policies only support vCenters or ESX hosts running vSphere 5.1 through 6.0.
- Note that VMware DRS cluster datastores are supported in VMware Copy Data and Use Data policies.
- 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 **50**.
- Note that VMware Copy Data policies do not support virtual machine SCSI controllers where the SCSI Bus Sharing value is set to virtual or physical.

- Note that Instant Access recoveries utilizing the VM Copy method are not supported at the datastore level. Instant Access 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 Copy Data policy.
- One or more schedules might also be associated with a policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

#### CONSIDERATIONS FOR VMWARE VIRTUAL VOLUMES:

- Storage Workflows that include virtual machines stored on virtual volume (VVOL) datastores through VM Copy sub-policies are supported. Replication is supported on the VM Copy 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 a VMware Copy Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, click **New** , then select **VMware**  in the Copy Data column. The VMware Copy Data Policy editor opens.
3. Click the **1: Source** tab. From the drop-down menu select **VMs and Templates** or **Storage**. From the list of available sites, select one or more resources to copy as defined by the associated Storage Workflows, including virtual machines, datastores, folders, vApps, and datacenters.
4. Click the **2. Storage Workflow** tab, then select a Storage Workflow that meets your copy data criteria.
5. Set a time to run the Storage Workflow with the **Specify Activation Time** trigger, then click **Add** to add it to the policy. To run the Storage Workflow on demand through the **Monitor**  tab, select **Manual**.

If configuring more than one Storage Workflow in a policy, use the **Same as workflow** option to trigger multiple Storage Workflows to run at the same time.

**Note:** Only Storage Workflows with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating a Storage Workflow.

6. Repeat the steps above to add additional Storage Workflows to the policy.
7. Click the **3. Options** tab. Set the policy 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 policy application or file system consistent, or making specific virtual machines included in the policy application or file system consistent. Application consistent copy data captures data in memory and transactions in process.

### VM Snapshot Scripts

VM snapshot prescripts and postscripts are scripts that can be run before or after a VMware virtual machine snapshot runs. The snapshot prescript runs before a VMware virtual machine snapshot is run, while the snapshot postscript runs after the snapshot completes. Select a virtual machine, then from the Scripts field enter individual script commands or the location of the prescript or postscript.

A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the `/data/userscripts` directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: `/data/userscripts/<name of script>`.

### Skip read only datastores

Enable to skip datastores mounted as read-only in Vcenter.

### Enable Policy-Level Scripts

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the `/data/userscripts` directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: `/data/userscripts/<name of script>`. For information about script return codes, see [Return Code Reference](#) on page 409.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

### Enable Policy-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot subpolicy runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. Enter individual script commands or the location of the prescript or postscript.

A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the `/data/userscripts` directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: `/data/userscripts/<name of script>`.

**\_SNAPSHOTS\_** is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the VMware virtual machine snapshots created by the policy. The format of each value is as follows: <virtual machine name>:<volume name>:<snapshot name>.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.


8. Click the **4: Notification** tab. Select the policy 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. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.

#### NEXT STEPS:

- If in a Linux environment, consider creating VADP proxies to enable load sharing. See [Create VMware Copy Data Policy Proxies](#) on page **175**.
- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page **235**.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Use the Catalog Browse feature to review the recovery point. See [Browse Catalog](#) on page **252**.
- Create a VMware Use Data policy. See [Create a VMware Use Data Policy](#) on page **213**.

#### RELATED TOPICS:

- [NetApp ONTAP Document: Cluster and Vserver Peering Express Guide](#)
- [Plan Overview](#) on page **99**
- [Edit a Policy](#) on page **230**
- [Delete a Policy](#) on page **231**
- [Create a Schedule](#) on page **119**

- [Create VMware Copy Data Policy Proxies](#) on page **175**
- [Create a VMware Use Data Policy](#) on page **213**

## Create VMware Copy Data Policy Proxies

In IBM Spectrum Copy Data Management, running VMware Copy Data policies through VADP can be taxing on your system resources. By creating VMware Copy Data policy proxies, you enable load sharing and load balancing for those policies in Linux environments.

Note that the first time a given policy is run, the proxies do not take effect because VM clone technology is used. But the second and subsequent times the policy 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 Copy Data 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 policies 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 policy.

### 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.

### 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/ECX**.
  - 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 **ecxvadb** is started on the proxy machine. A log file **ecxvadb.log** is generated in `/opt/ECX/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 Copy Data policy. 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 Copy Data policy. 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/ECX`):

```
./uninstall_vmdkbackup
```

The installation directory is removed.



**RELATED TOPICS:**

- [Create a VMware Copy Data Policy](#) on page **170**
- [Start, Stop, and Hold a Job Session](#) on page **235**

## Use Data Policies

[Create an Application Use Policy](#) on page **179**

[Create a DellEMC Unity Use Data Policy](#) on page **186**

[Create an IBM Spectrum Accelerate Use Data Policy](#) on page **191**

[Create an IBM Spectrum Virtualize Use Data Policy](#) on page **196**

[Create a NetApp ONTAP Use Data Policy](#) on page **201**

[Create a Pure Storage FlashArray Use Data Policy](#) on page **208**

[Create a VMware Use Data Policy](#) on page **213**

## Create an Application Use Policy

IBM Spectrum Copy Data Management leverages Copy Data Management technology for recovering application databases through Application Use policies. Your Oracle clones can be utilized and consumed instantly through IBM Spectrum Copy Data Management Instant Access policies. 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. Snapshot 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 Access or Instant Recovery to a new location using a masked image, created through an Application Copy policy with data masking enabled, **Instant Database Recovery**, which provides Instant Access or Instant Recovery using a non-masked image and point-in-time transaction logs, and **Instant Access**, which mounts a database for RMAN restores with application and operating system support.

The following Microsoft SQL Server workflows are supported: **Instant Database Recover**, which provides Instant Access or Instant Recovery using a non-masked image and point-in-time transaction logs, and **Instant Access**, which mounts a database for restores with application and operating system support. Databases can be recovered from a standalone instance to an AlwaysOn Availability Group, as well as from an AlwaysOn Availability Group to a standalone instance.

### BEFORE YOU BEGIN:

- Create and run an Application Copy policy. See [Create an Application Copy Policy](#) on page 146.
- Review Oracle requirements. See [Oracle Requirements](#) on page 31.
- 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.
- Note that an automated disk mount of an IBM Spectrum Copy Data Management copy on LVM storage must not have the original disk for that copy still present and mounted while the lvm2-lvmetad service is active on the system. If Oracle data resides on LVM volumes, you must stop and disable the lvm2-lvmetad service before running Application Copy or Use jobs. To disable the lvm2-lvmetad, 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 policy, add SMTP resources. See [Register a Provider](#) on page 50.
- You must add credentials to the destination virtual machine when recovering with the subnet option. See [Add Credentials to a Virtual Machine](#) on page 69.

#### ORACLE DATABASE CONSIDERATIONS:








- Note that RAC databases cannot be recovered to their original location while the protected (source) database is still running, even if a new name is chosen for the recovered database. You must either recover the database to another cluster, or shut down the source database before recovering a copy to the same cluster.
- 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 copy 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.






#### 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 Access or Instant Database Recovery policy, 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 policy.

**Best Practice:** Create a schedule before creating a policy so that you can easily add the schedule to the job definition in the New Policy editor.

#### **To create an Oracle Application Use policy:**

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **Application**  in the Use Data column. The Application Use Policy editor opens.
3. Click the **1: Workflow** tab. Select **Oracle**.
4. Select a template. Available options include **DevOps** , **Instant Database Recovery** , and **Instant Access** .

5. Click **Source**  . 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.
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**. If recovery from one snapshot fails, another copy from the same site is used.  
If creating an **Instant Database Recovery**  policy, 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 Access**  policy, 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 policy 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. Click the **2: Options** tab. Select the policy options.

### Make Permanent

Set the default permanent restoration action of the policy. 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 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 protocol is available, select the protocol to take priority in the policy. Available protocols include iSCSI and Fibre Channel.

### Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Use Data policy 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.

### Record mounted copies in RMAN local repository (Instant Access only)

Select this option to catalog mounted copies into RMAN at the end of the Instant Access job. This option is an alternative to the cataloging of IBM Spectrum Copy Data Management-created copies during a Copy Data job. If the RMAN cataloging option is selected in the copy policy, 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 Access job runs.

### Enable Policy-Level Scripts

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>. For information about script return codes, see [Return Code Reference](#) on page 409.

**\_VOLUMES\_** is an optional parameter for Use Data policy postscripts that displays a comma separated value string containing the storage volumes created by the Use Data policy. The format of each value is as follows: <registered provider name>:<volume name>.

For Use Data policy postscripts 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 222. **State** and **status** arguments are not supported for Copy Data policies.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

9. Click the **3: Notification** tab. Select the policy 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. Click the **4: Schedule** tab. Select **Start job now** to run the job immediately, as defined by the order of the recovery sources in the **1. Workflow** tab. 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 **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

11. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.

### To create a Microsoft SQL Application Use policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **Application**  in the Use Data column. The Application Use Policy editor opens.
3. Click the **1: Workflow** tab. Select **Microsoft SQL (Standalone)**  or **Microsoft SQL (Always On)** .
4. Select a template. Available options include **Instant Database Recovery**  for Microsoft SQL standalone and Always On policies and **Instant Access**  for Microsoft SQL Always On policies.
5. Click **Source** . 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.
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**. If recovery from one snapshot fails, another copy from the same site is used.  
If creating an **Instant Database Recovery**  policy, 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.
7. Click **Destination** . Select a source site and an associated Microsoft SQL database. Click the Destination field to enter an optional alternate name for the database.
8. Click the **2: Options** tab. Select the policy options.

#### Make Permanent

Set the default permanent restoration action of the policy. 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 protocol is available, select the protocol to take priority in the policy. Available protocols include iSCSI and Fibre Channel.

#### Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Use Data policy 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.

### **Overwrite existing database**

Enable to detach the original instance database and attach the new database upon recovery.

### **Enable Policy-Level Scripts**

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>. For information about script return codes, see [Return Code Reference](#) on page 409.

**\_VOLUMES\_** is an optional parameter for Use Data policy postscripts that displays a comma separated value string containing the storage volumes created by the Use Data policy. The format of each value is as follows: <registered provider name>:<volume name>.

For Use Data policy postscripts 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 222. **State** and **status** arguments are not supported for Copy Data policies.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

9. Click the **3: Notification** tab. Select the policy 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. Click the **4: Schedule** tab. Select **Start job now** to run the job immediately, as defined by the order of the recovery sources in the **1. Workflow** tab. 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 **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

11. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.

#### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page **235**.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- 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 Application Copy Policy](#) on page **146**
- [Plan Overview](#) on page **99**
- [Edit a Policy](#) on page **230**
- [Delete a Policy](#) on page **231**
- [Create a Schedule](#) on page **119**

## Create a DellEMC Unity Use Data Policy

IBM Spectrum Copy Data Management leverages Copy Data Management technology for recovering DellEMC Unity volumes through DellEMC Unity Use Data policies.

### Instant Access

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 Copy Data policy. 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 Copy Data policies. See [Create a DellEMC Unity Copy Data Policy](#) on page **154**.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page **50**.

#### 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 policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page **119**.

### *To create an Instant Access DellEMC Unity Use Data policy:*

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **DellEMC Unity**  in the Use Data column. The DellEMC Unity Use Data Policy editor opens.
3. Click the **1: Workflow** tab. Select the **Instant Access**  template.
4. 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.

5. 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**. If recovery from one snapshot fails, another copy from the same site is used.

**Note:** When selecting a specific version, data created through VMware Copy Data policies 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.

6. 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.

7. Click **2. Options**. Select the 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 Use Data policy 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.

#### **Enable Policy-Level Scripts**

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>. For information about script return codes, see [Return Code Reference](#) on page 409.

**\_VOLUMES\_** is an optional parameter for Use Data policy postscripts that displays a comma separated value string containing the storage volumes created by the Use Data policy. The format of each value is as follows: <registered provider name>:<volume name>.

For Use Data policy postscripts 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 **222**. **State** and **status** arguments are not supported for Copy Data policies.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.



8. Click the **3: Notification** tab. Select the policy 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. Click the **4: Schedule** tab. Select **Start job now** to run the job immediately, as defined by the order of the recovery sources in the **1: Workflow** tab. 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. A policy paired with a schedule is a job.  
**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.
10. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.

### To create a Restore Volume(s) DellEMC Unity Use Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **DellEMC Unity**  in the Use Data column. The DellEMC Unity Use Data Policy editor opens.
3. Click the **1: Workflow** tab. Select the **Restore Volume(s)**  template.
4. 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.
5. 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**. 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.

- Click **2. Options**. Select the 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 Use Data policy 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.

#### **Enable Policy-Level Scripts**

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>. For information about script return codes, see [Return Code Reference](#) on page 409.

**\_VOLUMES\_** is an optional parameter for Use Data policy postscripts that displays a comma separated value string containing the storage volumes created by the Use Data policy. The format of each value is as follows: <registered provider name>:<volume name>.

For Use Data policy postscripts 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 222. **State** and **status** arguments are not supported for Copy Data policies.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

- Click the **3: Notification** tab. Select the policy 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.

- Click the **4: Schedule** tab. Select **Start job now** to run the job immediately, as defined by the order of the recovery sources in the **1: Workflow** tab. 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

10. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.

#### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page **235**.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

#### RELATED TOPICS:

- [Plan Overview](#) on page **99**
- [Create a DellEMC Unity Copy Data Policy](#) on page **154**
- [Using State and Status Arguments in Postscripts](#) on page **222**
- [Edit a Policy](#) on page **230**
- [Delete a Policy](#) on page **231**
- [Create a Schedule](#) on page **119**
- [Search and Filter Guidelines](#) on page **404**

## Create an IBM Spectrum Accelerate Use Data Policy

IBM Spectrum Copy Data Management leverages Copy Data Management technology for recovering IBM Spectrum Accelerate volumes through IBM Spectrum Accelerate Use Data policies.

### Instant Access

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 Copy Data policy. 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 Copy Data policies. See [Create an IBM Spectrum Accelerate Copy Data Policy](#) on page 157.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page 50.

#### 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 policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

### **To create an Instant Access IBM Spectrum Accelerate Use Data policy:**

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **IBM Spectrum Accelerate**  in the Use Data column. The IBM Spectrum Accelerate Use Data Policy editor opens.

3. Click the **1: Workflow** tab. Select the **Instant Access**  template.
4. Click **Source** . 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.
5. 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**. If recovery from one snapshot fails, another copy from the same site is used.

**Note:** When selecting a specific version, data created through VMware Copy Data policies that apply to the selected IBM resource display, as the same data is contained with the snapshot for VMware and non-VMware related data.

6. 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.

7. Click **2. Options**. Select the 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 Use Data policy 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.

#### **Enable Policy-Level Scripts**

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>. For information about script return codes, see [Return Code Reference](#) on page 409.

**\_VOLUMES\_** is an optional parameter for Use Data policy postscripts that displays a comma separated value string containing the storage volumes created by the Use Data policy. The format of each value is as follows: <registered provider name>:<volume name>.



For Use Data policy postscripts 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 **222**. **State** and **status** arguments are not supported for Copy Data policies.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

8. Click the **3: Notification** tab. Select the policy 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. Click the **4: Schedule** tab. Select **Start job now** to run the job immediately, as defined by the order of the recovery sources in the **1: Workflow** tab. 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

10. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.

### **To create a Restore Volume(s) IBM Spectrum Accelerate Use Data policy:**

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **IBM Spectrum Accelerate**  in the Use Data column. The IBM Spectrum Accelerate Use Data Policy editor opens.
3. Click the **1: Workflow** tab. Select the **Restore Volume(s)**  template.
4. Click **Source** . 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.
5. 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**. If recovery from one snapshot fails, another copy from the same site is used.
6. 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.

7. Click **2. Options**. Select the 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 Use Data policy 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.

### Enable Policy-Level Scripts

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>. For information about script return codes, see [Return Code Reference](#) on page 409.

**\_VOLUMES\_** is an optional parameter for Use Data policy postscripts that displays a comma separated value string containing the storage volumes created by the Use Data policy. The format of each value is as follows: <registered provider name>:<volume name>.

For Use Data policy postscripts 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 222. **State** and **status** arguments are not supported for Copy Data policies.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

8. Click the **3: Notification** tab. Select the policy 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. Click the **4: Schedule** tab. Select **Start job now** to run the job immediately, as defined by the order of the recovery sources in the **1: Workflow** tab. 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

10. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.

**NEXT STEPS:**

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page **235**.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

**RELATED TOPICS:**

- [Plan Overview](#) on page **99**
- [Create an IBM Spectrum Accelerate Copy Data Policy](#) on page **157**
- [Using State and Status Arguments in Postscripts](#) on page **222**
- [Edit a Policy](#) on page **230**
- [Delete a Policy](#) on page **231**
- [Create a Schedule](#) on page **119**
- [Search and Filter Guidelines](#) on page **404**

## Create an IBM Spectrum Virtualize Use Data Policy

IBM Spectrum Copy Data Management leverages Copy Data Management technology for recovering IBM Spectrum Virtualize volumes through IBM Spectrum Virtualize Use Data policies.

### Instant Access

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 Copy Data policy. 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 Copy Data policies. See [Create an IBM Spectrum Virtualize Copy Data Policy](#) on page 160.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page 50.

#### 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 policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

### To create an Instant Access IBM Spectrum Virtualize Use Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **IBM Spectrum Virtualize**  in the Use Data column. The IBM Spectrum Virtualize Use Data Policy editor opens.

3. Click the **1: Workflow** tab. Select the **Instant Access**  template.
4. Click **Source** . 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.
5. 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**. If recovery from one snapshot fails, another copy from the same site is used.

**Note:** When selecting a specific version, data created through VMware Copy Data policies that apply to the selected IBM resource display, as the same data is contained with the snapshot for VMware and non-VMware related data.

6. 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.

7. Click **2. Options**. Select the policy 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 Use Data policy 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.

#### **Enable Policy-Level Scripts**

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>. For information about script return codes, see [Return Code Reference](#) on page 409.

**\_VOLUMES\_** is an optional parameter for Use Data policy postscripts that displays a comma separated value string containing the storage volumes created by the Use Data policy. The format of each value is as follows: <registered provider name>:<volume name>.

For Use Data policy postscripts 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 **222**. **State** and **status** arguments are not supported for Copy Data policies.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

- Click the **3: Notification** tab. Select the policy 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.


- Click the **4: Schedule** tab. Select **Start job now** to run the job immediately, as defined by the order of the recovery sources in the **1: Workflow** tab. 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

- Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.

### To create a Restore Volume(s) IBM Spectrum Virtualize Use Data policy:

- Click the **Plan**  tab. On the Views pane, select **Policies** .
- Click the **All Policies** tab, select **New** , then select **IBM Spectrum Virtualize**  in the Use Data column. The IBM Spectrum Virtualize Use Data Policy editor opens.
- Click the **1: Workflow** tab. Select the **Restore Volume(s)**  template.
- Click **Source** . 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.
- 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**. If recovery from one snapshot fails, another copy from the same site is used.

6. 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.

7. Click **2. Options**. Select the 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 Use Data policy 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.

#### **Enable Policy-Level Scripts**

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>. For information about script return codes, see [Return Code Reference](#) on page 409.

**\_VOLUMES\_** is an optional parameter for Use Data policy postscripts that displays a comma separated value string containing the storage volumes created by the Use Data policy. The format of each value is as follows: <registered provider name>:<volume name>.

For Use Data policy postscripts 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 222. **State** and **status** arguments are not supported for Copy Data policies.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

8. Click the **3: Notification** tab. Select the policy 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. Click the **4: Schedule** tab. Select **Start job now** to run the job immediately, as defined by the order of the recovery sources in the **1: Workflow** tab. 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

10. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.

#### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page **235**.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

#### RELATED TOPICS:

- [Plan Overview](#) on page **99**
- [Create an IBM Spectrum Virtualize Copy Data Policy](#) on page **160**
- [Using State and Status Arguments in Postscripts](#) on page **222**
- [Edit a Policy](#) on page **230**
- [Delete a Policy](#) on page **231**
- [Create a Schedule](#) on page **119**
- [Search and Filter Guidelines](#) on page **404**



## Create a NetApp ONTAP Use Data Policy

IBM Spectrum Copy Data Management leverages Copy Data Management technology for recovering NetApp ONTAP volumes and files through NetApp ONTAP Use Data policies.

### Instant Access

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 Copy Data policy. 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 policies 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 Copy Data policy. Files are restored to their original location.

#### BEFORE YOU BEGIN:








- Create and run one or more NetApp ONTAP Copy Data policies. See [Create a NetApp ONTAP Copy Data Policy](#) on page 164.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page 50.

#### CONSIDERATIONS:


- Note that a volume restore through a NetApp ONTAP Use Data policy that includes NetApp ONTAP storage systems operating in 7-mode is not supported. These providers will not display during policy creation.
- Note that a file restore from a mirror location is not available for NetApp ONTAP storage systems operating in 7-mode.
- Restore Volume policies 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 Use Data policy 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 Access Use Data policy 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 policies 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 Use Data policies 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 policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

### To create an Instant Access NetApp ONTAP Use Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **NetApp ONTAP**  in the Use Data column. The NetApp ONTAP Use Data Policy editor opens.
3. Click the **1: Workflow** tab. Select the **Instant Access**  template.
4. Click **Source** . 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.
5. 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**. If recovery from one snapshot fails, another copy from the same site is used.

**Note:** When selecting a specific version, data created through VMware Copy Data policies 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.

6. 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 Access volume is made permanent, you must rename the NFS/CIFS path for the permanent volume prior to running the same job again.

7. Click **2. Options**. Select the policy 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 Use Data policy 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.

#### **Enable Policy-Level Scripts**

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>. For information about script return codes, see [Return Code Reference](#) on page 409.

**\_VOLUMES\_** is an optional parameter for Use Data policy postscripts that displays a comma separated value string containing the storage volumes created by the Use Data policy. The format of each value is as follows: <registered provider name>:<volume name>.

For Use Data policy postscripts 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 222. **State** and **status** arguments are not supported for Copy Data policies.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

8. Click the **3: Notification** tab. Select the policy 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. Click the **4: Schedule** tab. Select **Start job now** to run the job immediately, as defined by the order of the recovery sources in the **1: Workflow** tab. 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. A policy paired with a schedule is a job.
 

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.
10. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.

### To create a Restore Volume(s) NetApp ONTAP Use Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **NetApp ONTAP**  in the Use Data column. The NetApp ONTAP Use Data Policy editor opens.
3. Click the **1: Workflow** tab. Select the **Restore Volume(s)**  template.
4. Click **Source** . 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.
5. 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**. If recovery from one snapshot fails, another copy from the same site is used.
6. 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.
7. Click **2. Options**. Select the 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 Use Data policy 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.

### Enable Policy-Level Scripts

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>. For information about script return codes, see [Return Code Reference](#) on page 409.

**\_VOLUMES\_** is an optional parameter for Use Data policy postscripts that displays a comma separated value string containing the storage volumes created by the Use Data policy. The format of each value is as follows: <registered provider name>:<volume name>.

For Use Data policy postscripts 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 222. **State** and **status** arguments are not supported for Copy Data policies.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.



8. Click the **3: Notification** tab. Select the policy 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. Click the **4: Schedule** tab. Select **Start job now** to run the job immediately, as defined by the order of the recovery sources in the **1: Workflow** tab. 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. A policy paired with a schedule is a job.  
**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.
10. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.

### To create a Restore File(s) NetApp ONTAP Use Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .

2. Click the **All Policies** tab, select **New** , then select **NetApp ONTAP**  in the Use Data column. The NetApp ONTAP Use Data Policy editor opens.
3. Click the **1: Workflow** tab. Select the **Restore File(s)**  template.
4. 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.
5. 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.
6. 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.
7. Click **2. Options**. Select the 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 Use Data policy 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.

#### Enable Policy-Level Scripts

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>. For information about script return codes, see [Return Code Reference](#) on page 409.

**\_VOLUMES\_** is an optional parameter for Use Data policy postscripts that displays a comma separated value string containing the storage volumes created by the Use Data policy. The format of each value is as follows: <registered provider name>:<volume name>.

For Use Data policy postscripts 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 222. **State** and **status** arguments are not supported for Copy Data policies.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.



8. Click the **3: Notification** tab. Select the policy 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.

- Click the **4: Schedule** tab. Select **Start job now** to run 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. A policy paired with a schedule is a job.  
**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.
- Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.

### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page **235**.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

### RELATED TOPICS:

- [Plan Overview](#) on page **99**
- [Create a NetApp ONTAP Copy Data Policy](#) on page **164**
- [Using State and Status Arguments in Postscripts](#) on page **222**
- [Edit a Policy](#) on page **230**
- [Delete a Policy](#) on page **231**
- [Create a Schedule](#) on page **119**
- [Search and Filter Guidelines](#) on page **404**

## Create a Pure Storage FlashArray Use Data Policy

IBM Spectrum Copy Data Management leverages Copy Data Management technology for recovering Pure Storage FlashArray volumes through Pure Storage FlashArray Use Data policies.

### Instant Access

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 Copy Data policy. Volumes can be restored to their original location or a new volume in the same or different Pure Storage system.









#### BEFORE YOU BEGIN:

- Create and run one or more Pure Storage FlashArray Copy Data policies. See [Create a Pure Storage FlashArray Copy Data Policy](#) on page 167.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page 50.

#### CONSIDERATIONS:

- One or more schedules might also be associated with a policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

### To create an Instant Access Pure Storage FlashArray Use Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **Pure Storage FlashArray**  in the Use Data column. The Pure Storage FlashArray Use Data Policy editor opens.
3. Click the **1: Workflow** tab. Select the **Instant Access**  template.
4. Click **Source** . 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.
5. 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**. If recovery from one snapshot fails, another copy from the same site is used.
6. Click **Destination** . Select a Pure Storage FlashArray destination.
7. Click **2. Options**. Select the 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 Use Data policy 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.

**Enable Policy-Level Scripts**

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the `/data/userscripts` directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: `/data/userscripts/<name of script>`. For information about script return codes, see [Return Code Reference](#) on page 409.

`_VOLUMES_` is an optional parameter for Use Data policy postscripts that displays a comma separated value string containing the storage volumes created by the Use Data policy. The format of each value is as follows: `<registered provider name>:<volume name>`.

For Use Data policy postscripts 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 222. **State** and **status** arguments are not supported for Copy Data policies.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

8. Click the **3: Notification** tab. Select the policy 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. Click the **4: Schedule** tab. Select **Start job now** to run the job immediately, as defined by the order of the recovery sources in the **1: Workflow** tab. 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

- Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.

### To create a Restore Volume(s) Pure Storage FlashArray Use Data policy:

- Click the **Plan**  tab. On the Views pane, select **Policies** .
- Click the **All Policies** tab, select **New** , then select **Pure Storage FlashArray**  in the Use Data column. The Pure Storage FlashArray Use Data Policy editor opens.
- Click the **1: Workflow** tab. Select the **Restore Volume(s)**  template.
- Click **Source** . 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.
- 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**. 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.
- Click **2. Options**. Select the 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 Use Data policy 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.

#### Enable Policy-Level Scripts

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>. For information about script return codes, see [Return Code Reference](#) on page 409.

**\_VOLUMES\_** is an optional parameter for Use Data policy postscripts that displays a comma separated value string containing the storage volumes created by the Use Data policy. The format of each value is as follows: <registered provider name>:<volume name>.

For Use Data policy postscripts 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 **222**. **State** and **status** arguments are not supported for Copy Data policies.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

8. Click the **3: Notification** tab. Select the policy 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. Click the **4: Schedule** tab. Select **Start job now** to run the job immediately, as defined by the order of the recovery sources in the **1: Workflow** tab. 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

10. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.

#### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page **235**.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

#### RELATED TOPICS:

- [Plan Overview](#) on page **99**
- [Create a Pure Storage FlashArray Copy Data Policy](#) on page **167**

- [Using State and Status Arguments in Postscripts](#) on page **222**
- [Edit a Policy](#) on page **230**
- [Delete a Policy](#) on page **231**
- [Create a Schedule](#) on page **119**
- [Search and Filter Guidelines](#) on page **404**

## Create a VMware Use Data Policy

IBM Spectrum Copy Data Management leverages Copy Data Management technology for testing and cloning use cases, instant recovery, and full disaster recovery. VMware Copy Data policies support Instant Virtualization and Instant Access scenarios.

**Instant Virtualization**  policies 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 policies 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 policies 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 Access

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 Copy Data policy. See [Create a VMware Copy Data Policy](#) on page 170.

- 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 policy, add SMTP resources. See [Register a Provider](#) on page 50.
- You must add credentials to the destination virtual machine when recovering with the subnet option. See [Add Credentials to a Virtual Machine](#) on page 69.

**CONSIDERATIONS:**

- Note that VMware Copy Data and Use Data policies only support vCenters or ESX hosts running vSphere 5.1 through 6.0.
- If a recovery through an Instant Virtualization RRP job using an AWS-based Copy Data Workflow 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.
- Note that VMware DRS cluster datastores are supported in VMware Copy Data and Use Data policies.
- Note that after a Use Data policy through Instant Access policy 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 Access recoveries utilizing the VM Copy method are not supported at the datastore level. Instant Access datastore level recoveries are supported through the primary storage snapshot method.
- Note that in Instant Virtualization 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 Access Use Data recoveries of VMDKs through snapshots of IBM Spectrum Protect Snapshot-protected virtual machines are not supported.
- Instant Virtualization Use Data 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 Use Data policy 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 Use Data policies 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 policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.




#### CONSIDERATIONS FOR VMWARE VIRTUAL VOLUMES:

- All functionality of Use Data 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 Use Data workflows, as IBM Spectrum Copy Data Management stores the VM Copy 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 policy so that you can easily add the schedule to the job definition in the New Policy editor.

#### To create an Instant Virtualization VMware Use Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **VMware**  in the Use Data column. The VMware Use Data Policy editor opens.
3. Click the **1: Workflow** tab. Select the **Instant Virtualization**  template.

4. Click **Source**  . Select a source site and an associated VMware source to view virtual machines, 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.
5. 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**. If recovery from one snapshot fails, another copy from the same site is used.
6. Click **Destination**  . Select a source site and an associated 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**.
7. Select the virtual network and datastores mapping options if you selected a destination different from the original host or cluster. The Virtual Networks pane displays all of the virtual networks associated with your VMware Use Data policy 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 Virtualization 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 original subnets and IP addresses for VM guest OS on destination** option is enabled. Disable this option to create a new mapping by selecting **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. See [Add Credentials to a Virtual Machine](#) on page 69.

8. Click the **2: Options** tab. Select the policy options.

### Default Mode

Set the VMware Use Data policy to run in Test, Production, or Clone mode by default. Once the policy is created, it can be run in Test, Production, or Clone mode through the **Monitor**  tab.

### Protocol Priority

If more than one storage protocol is available, select the protocol to take priority in the policy. 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.

### Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Use Data policy 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.

### Enable Policy-Level Scripts

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>. For information about script return codes, see [Return Code Reference](#) on page 409.

**\_VOLUMES\_** is an optional parameter for Use Data policy postscripts that displays a comma separated value string containing the storage volumes created by the Use Data policy. The format of each value is as follows: <registered provider name>:<volume name>.

For Use Data policy postscripts 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 **222**. **State** and **status** arguments are not supported for Copy Data policies.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

9. Click the **3: Notification** tab. Select the policy 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. Click the **4: Schedule** tab. Select **Start job now** to run the job immediately, as defined by the order of the recovery sources in the **1: Workflow** tab. 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 **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.





11. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.
12. Once the policy completes successfully, select one of the following options from the **Actions** menu on the General tab of the job session on the Monitor 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 policy Options 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 policy Options screen. This option migrates the virtual machine through vMotion to the Datastore and Virtual Network defined as the "For Test" network.

### To create an Instant Access VMware Use Data policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **VMware**  in the Use Data column. The VMware Use Data Policy editor opens.

3. Click the **1: Workflow** tab. Select the **Instant Access**  template.
4. 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, 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.
5. 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**. If recovery from one snapshot fails, another copy from the same site is used.
6. 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**.
7. Select the datastore and virtual disk mapping options if you selected a destination different from the original host or cluster.

### Virtual Disks

Select virtual machine destinations. 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.

### Datastores

Set the destination datastore.

8. Click the **2: Options** tab. Select the policy options.

### Protocol Priority

If more than one storage protocol is available, select the protocol to take priority in the policy. 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 Use Data policy 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.

### Enable Policy-Level Scripts

Policy-level prescripts and postscripts are scripts that can be run before or after a policy runs. Enter individual script commands or the location of the prescript or postscript.

Policy-level scripts are run at the policy-level. A script can consist of one or many commands, such as a shell script. Note that external scripts must be added to the /data/userscripts directory on the IBM Spectrum Copy Data Management appliance. To invoke an external script through the Prescript and Postscript fields, enter the full path and the name of the script: /data/userscripts/<name of script>. For information about script return codes, see [Return Code Reference](#) on page 409.

**\_VOLUMES\_** is an optional parameter for Use Data policy postscripts that displays a comma separated value string containing the storage volumes created by the Use Data policy. The format of each value is as follows: <registered provider name>:<volume name>.

For Use Data policy postscripts 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 222. **State** and **status** arguments are not supported for Copy Data policies.

Select **Continue operation on script failure** to continue running the policy if a command in any of the scripts associated with the policy fails.

9. Click the **3: Notification** tab. Select the policy 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. Click the **4: Schedule** tab. Select **Start job now** to run the job immediately, as defined by the order of the recovery sources in the **1: Workflow** tab. 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 **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

11. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**. The policy runs as defined by your triggers, or can be run manually from the **Monitor**  tab.

### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page **235**.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- 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 VMware Copy Data Policy](#) on page **170**
- [Using State and Status Arguments in Postscripts](#) on page **222**
- [NetApp ONTAP Document: Cluster and Vserver Peering Express Guide](#)
- [NetApp ONTAP Document: iSCSI Configuration and Provisioning for ESX Express Guide](#)
- [Plan Overview](#) on page **99**
- [Edit a Policy](#) on page **230**
- [Delete a Policy](#) on page **231**
- [Create a Schedule](#) on page **119**
- [Search and Filter Guidelines](#) on page **404**

## Using State and Status Arguments in Postscripts

The functionality described in this topic (**state** and **status** arguments) applies only to postscripts for Use Data policies. Copy Data policies are not supported. Prescripts are not supported.

Certain positional arguments can be passed to a Use Data policy 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 Use Data 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 **Postscript** field:

```
/data/userscripts/<name of script> _STATE_ _STATUS_
```

The following is a simple example of code that can be added to a postscript (or can be its own postscript) in order to see the specific state and status values in the job log:

```
echo state $1
echo status $2
```

In the above script, \$1 and \$2 represent the arguments that are being passed in to the script from the **Postscript** field in the IBM Spectrum Copy Data Management user interface.

So for the above postscript example, the following output would display if the job ran in an Instant Access 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 policy type (such as Copy Data policy scripts), simply passes “`_STATE_`” or “`_STATUS_`” as plain text.

The valid values for state and status vary depending on the Use Data policy type. Following is a table of valid state and status values.

Valid **state** and **status** values

Use Data policy type	State	Status
DellEMC Unity– Instant Access	<ul style="list-style-type: none"> <li>• IA</li> <li>• IA_END</li> <li>• IA_CANCEL</li> </ul>	<ul style="list-style-type: none"> <li>• SUCCESS</li> <li>• FAILED</li> <li>• PARTIAL</li> </ul>
DellEMC Unity– Restore Volume(s)	<ul style="list-style-type: none"> <li>• RESTORE_VOLUME</li> <li>• RESTORE_VOLUME_CANCEL</li> </ul>	<ul style="list-style-type: none"> <li>• SUCCESS</li> <li>• FAILED</li> <li>• PARTIAL</li> </ul>
IBM – Instant Access	<ul style="list-style-type: none"> <li>• IA</li> <li>• IA_END</li> <li>• IA_CANCEL</li> <li>• IA_PERMANENT</li> </ul>	<ul style="list-style-type: none"> <li>• SUCCESS</li> <li>• FAILED</li> <li>• PARTIAL</li> </ul>
IBM – Restore Volume (s)	<ul style="list-style-type: none"> <li>• RESTORE_VOLUME</li> <li>• RESTORE_VOLUME_CANCEL</li> </ul>	<ul style="list-style-type: none"> <li>• SUCCESS</li> <li>• FAILED</li> <li>• PARTIAL</li> </ul>
NetApp ONTAP– Instant Access	<ul style="list-style-type: none"> <li>• IA</li> <li>• IA_END</li> <li>• IA_CANCEL</li> <li>• IA_PERMANENT</li> </ul>	<ul style="list-style-type: none"> <li>• SUCCESS</li> <li>• FAILED</li> <li>• PARTIAL</li> </ul>
NetApp ONTAP– Restore Volume(s)	<ul style="list-style-type: none"> <li>• RESTORE_VOLUME</li> <li>• RESTORE_VOLUME_CANCEL</li> </ul>	<ul style="list-style-type: none"> <li>• SUCCESS</li> <li>• FAILED</li> <li>• PARTIAL</li> </ul>
NetApp ONTAP– Restore Files	<ul style="list-style-type: none"> <li>• RESTORE_FILE</li> <li>• RESTORE_FILE_CANCEL</li> </ul>	<ul style="list-style-type: none"> <li>• SUCCESS</li> <li>• FAILED</li> <li>• PARTIAL</li> </ul>
VMware – Instant Virtualization	<ul style="list-style-type: none"> <li>• IV_TEST</li> <li>• IV_TEST_END</li> <li>• IV_CLONE</li> <li>• IV_PRODUCTION</li> </ul>	<ul style="list-style-type: none"> <li>• SUCCESS</li> <li>• FAILED</li> <li>• PARTIAL</li> </ul>

Use Data policy type	State	Status
	<ul style="list-style-type: none"><li>• IV_CANCEL</li></ul>	
VMware – Instant Access	<ul style="list-style-type: none"><li>• IA</li><li>• IA_END</li><li>• IA_CANCEL</li><li>• IA_PERMANENT</li></ul>	<ul style="list-style-type: none"><li>• SUCCESS</li><li>• FAILED</li><li>• PARTIAL</li></ul>

**RELATED TOPICS:**

- [Create a DellEMC Unity Use Data Policy](#) on page **186**
- [Create an IBM Spectrum Virtualize Use Data Policy](#) on page **196**
- [Create a NetApp ONTAP Use Data Policy](#) on page **201**
- [Create a VMware Use Data Policy](#) on page **213**



## Create a Script Policy

A Script policy defines a set of commands to run on the IBM Spectrum Copy Data Management appliance. Use the Script policy feature to add functionality to IBM Spectrum Copy Data Management.

A Script policy includes a user-defined script as well as other policy 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 as a policy 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 119 and [Register a Provider](#) on page 50.

#### **To construct a Script policy:**

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, click **New** , then select **Script** in the System column. The Script Policy Editor opens.
3. Click the **1: Scripts** tab. On the Scripts pane, 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** .
4. Continue to add commands. Use the reorder, edit, and delete options to assist you. Your script can consist of one or many commands.
5. Click the **2: Options** tab after all the commands are entered. Specify the applicable options for running the policy.

### Run scripts in order

Select **Run scripts in order** for the commands in the policy to run sequentially. The policy 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 policy to stop running as soon as one of the commands fails. Clear this option for the policy to continue to run after one of the commands fails. This option is only applicable if **Run scripts in order** is selected.

6. Click the **3: Notification** tab. Select the notification options for your policy.


### 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.

7. Click the **4: Schedule** tab. Select **Start job now** to create a policy 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

8. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**.

**Note:** If you selected the **Start job now** option, the job runs.

9. Click the **All Policies** tab. Your new policy appears in the Script Policy list.

### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page **235**.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

**RELATED TOPICS:**

- [Plan Overview](#) on page **99**
- [Return Code Reference](#) on page **409**
- [Edit a Policy](#) on page **230**
- [Delete a Policy](#) on page **231**
- [Create a Schedule](#) on page **119**





## Create a Report Policy

A Report policy 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 policy 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 policy. See [Create a Customized Report](#) on page 259.
- At least one provider must be associated with a Catalog Data policy. Before defining a catalog policy, add providers. See [Register a Provider](#) on page 50.
- For email notifications, at least one SMTP server must be configured. Before defining a policy, add SMTP resources. See [Register a Provider](#) on page 50.
- One or more schedules might also be associated with a policy. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#) on page 119.

### To create a Report policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** .
2. Click the **All Policies** tab, select **New** , then select **Report Policy**  in the Analyze Data column. The Report Policy editor opens.
3. Click the **1: Reports** tab. From the list of available reports, select one or more reports to include in the policy. Expand reports to view associated customized reports. Select **Default** to run the predefined report parameters.
4. Click the **2: Options** tab. Select render options for your emailed report attachments. Reports can be rendered as Adobe PDFs, Microsoft Word files, and Microsoft Excel spreadsheets.
5. Click the **3: Notification** tab. Select the notification options for your policy.

#### 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.

6. Click the **4: Schedule** tab. Select **Start job now** to create a policy 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. A policy paired with a schedule is a job.

**Note:** To create and select a new schedule, click **Views**, then select **Schedules** . Create a schedule, then return to the policy editor, refresh the Available Schedules pane, and select the new schedule.

7. Click the **5: Finish** tab. Enter a name for your policy and a meaningful description. When you are satisfied that the policy-specific information is correct, click **Finish**.

**Note:** If you selected the **Start job now** option, the job runs.

8. Click the **All Policies** tab. Your new policy appears in the policy list.

#### NEXT STEPS:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Stop, and Hold a Job Session](#) on page 235.
- Track the progress of the job session on the Monitor tab. See [Monitor a Job Session](#) on page 237.
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

#### RELATED TOPICS:

- [Plan Overview](#) on page 99
- [Report Overview](#) on page 256
- [Edit a Policy](#) on page 230
- [Delete a Policy](#) on page 231
- [Create a Schedule](#) on page 119

## Edit a Policy




Revise policy properties to change the provider that the policy is run against, Storage Workflows associated with a Copy Data policy, the policy options, notification and schedule properties.

**Note:** Copy Data policies created in IBM Spectrum Copy Data Management 2.1 are editable in IBM Spectrum Copy Data Management 2.2. Available editable options reflect the options that were available in IBM Spectrum Copy Data Management 2.1.

### BEFORE YOU BEGIN:

- Review the properties of your current policies. See [Plan Overview](#) on page **99**.

### To edit the properties of a policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** . Then click the **All Policies** tab.
2. Select the policy to edit by clicking in the row containing the policy name.
3. Click **Edit** . The Policy Editor opens.
4. Click through the Policy Editor Wizard making revisions as needed.
5. Click **Finish** when you are satisfied that the policy-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, Stop, and Hold a Job Session](#) on page **235**.
- Track the progress of the job on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- If SMTP options are enabled, an email message with information about the status of each task is sent when the job completes.

### RELATED TOPICS:


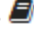

- [Plan Overview](#) on page **99**
- [Delete a Policy](#) on page **231**

## Delete a Policy

Delete a policy when it becomes obsolete. This keeps your operations current.

**Note:** The Maintenance policy removes providers and associated objects created by IBM Spectrum Copy Data Management when a Copy Data policy or Use Data policy in a pending state is deleted. This policy is run automatically.

### To delete a policy:

1. Click the **Plan**  tab. On the Views pane, select **Policies** . Then click the **All Policies** tab.
2. Select the policy to delete by clicking in the row containing the policy name.
3. Click **Delete** . A confirmation dialog box opens.
4. Confirm deletion. The policy is deleted.

#### RELATED TOPICS:

- [Plan Overview](#) on page **99**
- [Edit a Policy](#) on page **230**
- [Maintenance Policy](#) on page **232**

## Maintenance Policy

The Maintenance policy removes resources and associated objects created by IBM Spectrum Copy Data Management when a policy 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 policy 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 policy cannot be deleted.

The Maintenance policy only performs cleanup operations once a policy in a pending state is deleted. All logs associated with the deleted policy are removed from IBM Spectrum Copy Data Management, so it is advised to download job logs before the Maintenance policy's next run. The policy can be stopped and resumed; all pending operations set to occur before the policy was stopped will resume upon the next policy run.

After deleting a pending Application, DellEMC Unity, IBM, NetApp ONTAP, Pure Storage, or VMware Copy Data policy, all associated copy data, including recovery points, are deleted. The Maintenance policy removes all VM Copies and Primary copies associated with deleted VMware Copy Data and Use Data policies. Similarly, after deleting a pending DellEMC Unity, IBM, NetApp ONTAP, or Pure Storage Copy Data or Use Data policy, all associated DellEMC Unity, IBM, NetApp ONTAP, and Pure Storage locations are removed by the Maintenance policy. Once the Maintenance job completes, application, DellEMC Unity, IBM, NetApp ONTAP, Pure Storage or VMware data that was copied as part of the copy job cannot be recovered. Any data related to the deleted policy will not be recoverable.

The Maintenance policy also removes cataloged data associated with deleted Application, DellEMC Unity, IBM, NetApp ONTAP, Pure Storage and VMware Catalog Data policies, and removes jobs and job sessions related to Script and Report policies from the IBM Spectrum Copy Data Management interface.

### RELATED TOPICS:

- [Plan Overview](#) on page **99**
- [Delete a Policy](#) on page **231**





# Monitor

[Monitor Overview](#) on page **234**

[Start, Stop, and Hold a Job Session](#) on page **235**

[Monitor a Job Session](#) on page **237**

## Monitor Overview

Start, monitor, stop, and resume a job from the **Jobs**  pane on the **Monitor**  tab. From this pane, you can also view all scheduled and unscheduled job sessions, and start jobs before scheduled run times.



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. This information is found in the General tab.

Select the Jobs History tab in the job details window to view the duration and completion status from the prior runs of the selected job.

### RELATED TOPICS:

- [Start, Stop, and Hold a Job Session](#) on page **235**
- [Monitor a Job Session](#) on page **237**




## Start, Stop, and Hold a Job Session

From the **Jobs**  pane on the **Monitor**  tab you can run a job session on demand, stop a running job, and hold all future scheduled instances of a job from running until you are ready for the job to proceed.




### BEFORE YOU BEGIN:

- Create Catalog, script, and report policies. See [Plan Overview](#) on page 99.



### To start a job session:

1. Click the **Monitor**  tab. On the Views pane, select **Jobs** .
2. Select the job to run by clicking in the row containing the job name.
3. Click **Start** . A confirmation dialog box opens.  
**Note:** If a job session has multiple run options, such as running a Use Data policy in Test, Recovery, or Clone mode, you will be prompted to select a job session type.
4. Click **Yes**. The job session runs.
5. 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 stop a job session:

1. Click the **Monitor**  tab. On the Views pane, select **Jobs History** .
2. Select a running job to stop by clicking in the row containing the job name.
3. Click **Stop** . A confirmation dialog box opens.
4. Click **Yes**. The job session stops.

### To hold and release a job session:

1. Click the **Monitor**  tab. On the Views pane, select **Jobs** .
2. Select the job to suspend by clicking in the row containing the job name.
3. Click **Actions > Hold** to hold the job session. The job session status changes to **Held**, and all future scheduled instances of the job will not run until released.
4. Click **Actions > Release** to release the job session.

### NEXT STEPS:

- Track the progress of the job on the Monitor tab. See [Monitor a Job Session](#) on page **237**.
- Once the job session completes, review cataloged data through the Search and Report tabs. See [Browse Catalog](#) on page **252** and [Report Overview](#) on page **256**.

**RELATED TOPICS:**

- [Create a Schedule](#) on page **119**



## Monitor a Job Session

You can view the details of a job session that is currently running or one that has finished. Use the Monitor window 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, Stop, and Hold a Job Session](#) on page 235.



### To monitor a running job session:

1. Click the **Monitor**  tab. On the **Views** pane, select **Jobs** . The All Jobs pane 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.
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 **Monitor**  tab. On the **Views** pane, select **Jobs**  or **Jobs History** .
2. Click the drop-down arrow in the header of the **Type**, **Status**, or **Last Run Status** columns.
3. Select **Filters**, then choose a filter criteria.

### To view information about specific job sessions:

1. Click the **Monitor**  tab. In the **Views** pane, select **Jobs History** . The Jobs History pane displays job sessions including start and end date and time, policy name, job name, and status.
2. Select a job to view by clicking its job name. The following tabs open:

**General**

Displays an overview 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.

**Tasks**

Displays a task-by-task view of the job session, including start and end times, duration, and status.



**Log**

Displays the job log, which can be used for troubleshooting purposes. Job logs are automatically removed from the IBM Spectrum Copy Data Management Catalog six months after they are generated.

**Policy**

Displays an overview of the policy definition, including the policy name, sources, options, and notification settings.

**To see a list of jobs scheduled to run on a given day:**

1. Click the **Monitor**  tab. On the **Views** pane, select **Schedule** .
2. Select a monthly, weekly, or daily calendar view. A list of all jobs scheduled to run displays.

**RELATED TOPICS:**

- [Start, Stop, and Hold a Job Session](#) on page **235**
- [Collect Logs For Troubleshooting](#) on page **391**

# Search

[Search Overview](#) on page **240**

[Search for Objects](#) on page **241**

[View Object Details](#) on page **246**

[View NetApp ONTAP File Details](#) on page **247**

[Find and Restore a File](#) on page **249**

[Download Search Results](#) on page **251**

[Browse Catalog](#) on page **252**

## 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 Catalog may include CIFS shares, file systems, hosts, host containers, LUNs, NAS servers, NFS shares, pools, snapshots, storage resources, and storage resource replications.

The IBM Catalog may include FlashCopies, Hosts, IOGroups, MDisks, mirrors, Node Canisters, PortIPs, and volumes.

The NetApp ONTAP Catalog may include aggregates, CIFS shares, files, LUNs, networks, NFS exports, nodes, policies, protocols, qtrees, quotas, SnapMirrors, Snapshots, SnapVaults, SVMs, vFilers, and volumes.

The Recovery Catalog may contain datacenters, data stores, ESX hosts, LUNs, folders, recovery points, vApps, vDisks, vSnapshots, and vSpheres.

The VMware Catalog may include data stores, ESX hosts, LUNs, virtual disks, virtual machines, VMware hosts, and virtual snapshots.

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 Catalog Browser to browse through the list of providers. Drill into the Catalog 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 Catalog 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 241
- [View Object Details](#) on page 246
- [View NetApp ONTAP File Details](#) on page 247
- [Find and Restore a File](#) on page 249
- [Download Search Results](#) on page 251
- [Browse Catalog](#) on page 252



## 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 Catalog 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 50 and [Plan Overview](#) on page 99.

### 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 **251**.
- 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 **246**.

**RELATED TOPICS:**

- [Download Search Results](#) on page **251**
- [Browse Catalog](#) on page **252**
- [View Object Details](#) on page **246**
- [View NetApp ONTAP File Details](#) on page **247**
- [Search and Filter Guidelines](#) on page **404**
- [Select, Sort, and Reorder Columns](#) on page **407**

## 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 Catalog, 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 Catalog Data policy. 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. See [Create a NetApp ONTAP Storage Catalog Data Policy](#) on page **135**, [Create a VMware Catalog Data Policy](#) on page **143**, [Create an IBM Spectrum Virtualize Catalog Data Policy](#) on page **131**, [Create an IBM Spectrum Protect Snapshot Catalog Data Policy](#) on page **133**, and [Create a DellEMC Unity Catalog Data Policy](#) on page **127**.
- 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 Catalog, 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 **247**
- [Search Overview](#) on page **240**
- [Search for Objects](#) on page **241**

## 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 Catalog. 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 Catalog Data policy.

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 Catalog, 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 Catalog Data policy and a NetApp ONTAP File Catalog Data policy. By enabling the **Catalog all available snapshots** option in the NetApp ONTAP File Catalog Data policy, you can view multiple versions of a file. See [Create a NetApp ONTAP Storage Catalog Data Policy](#) on page 135 and [Create a NetApp ONTAP File Catalog Data Policy](#) on page 137.
- Search for a file by selecting **NetApp > File** from the Search For filter on the Advanced Search dialog. See [Search for Objects](#) on page 241.

### 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 Catalog, 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 **249**.

**RELATED TOPICS:**

- [View Object Details](#) on page **246**
- [Search Overview](#) on page **240**
- [Search for Objects](#) on page **241**
- [Find and Restore a File](#) on page **249**



## 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 a NetApp ONTAP File Catalog Data Policy](#) on page 137.
- 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 Policy 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 **247**
- [Search Overview](#) on page **240**
- [Search for Objects](#) on page **241**

## 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 241.

### 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 241
- [View Object Details](#) on page 246
- [View NetApp ONTAP File Details](#) on page 247

## Browse Catalog

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 Catalog may include CIFS shares, file systems, hosts, host containers, LUNs, NAS servers, NFS shares, pools, snapshots, storage resources, and storage resource replications.

The IBM Catalog may include FlashCopies, Hosts, IOGroups, MDisks, mirrors, Node Canisters, PortIPs, and volumes.

The NetApp ONTAP Catalog may include aggregates, CIFS shares, files, LUNs, networks, NFS exports, nodes, policies, protocols, qtrees, quotas, SnapMirrors, Snapshots, SnapVaults, SVMs, vFilers, and volumes.



The Recovery Catalog may contain datacenters, data stores, ESX hosts, LUNs, folders, recovery points, vApps, vDisks, vSnapshots, and vSpheres.

The VMware Catalog may include data stores, ESX hosts, LUNs, virtual disks, virtual machines, VMware hosts, and virtual snapshots.


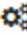
### BEFORE YOU BEGIN:

- You can only browse objects on providers that are registered and cataloged. See [Create a NetApp ONTAP Storage Catalog Data Policy](#) on page **135**, [Create a VMware Catalog Data Policy](#) on page **143**, [Create an IBM Spectrum Virtualize Catalog Data Policy](#) on page **131**, [Create an IBM Spectrum Protect Snapshot Catalog Data Policy](#) on page **133**, and [Create a DellEMC Unity Catalog Data Policy](#) on page **127**.

### To browse providers that are cataloged:

1. Click the **Search**  tab. On the Views pane, select **Catalog** . The Catalog Browser opens.
2. Drill down through providers in the Catalog 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 Catalog on a previous date:

1. Click the **Search**  tab. On the Views pane, select **Catalog** . The Catalog 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 Catalog Browser reflects the state of the Catalog on that date.
4. Drill down through providers in the Catalog 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 **240**
- [Search for Objects](#) on page **241**
- [View Object Details](#) on page **246**
- [View NetApp ONTAP File Details](#) on page **247**

# Report

[Report Overview](#) on page **256**

[Run a Report](#) on page **257**

[Create a Customized Report](#) on page **259**

[Download a Report](#) on page **261**

[Delete a Generated Report](#) on page **262**

[Application Reports](#) on page **263**

- [Application Configuration Report](#) on page **264**
- [Application RPO Compliance Report](#) on page **266**

[System Management Reports](#) on page **270**

- [Catalog Summary Report](#) on page **271**
- [Configuration Report](#) on page **273**
- [Policy Report](#) on page **277**
- [System Sizing Report](#) on page **279**

[File Analytics Reports](#) on page **281**

- [File Usage by Owner Report](#) on page **283**
- [Files By Age Report](#) on page **285**
- [Files By Category Report](#) on page **287**
- [Files By Size Report](#) on page **289**

[Protection Compliance Reports](#) on page **291**

- [DellEMC Unity RPO Compliance Report](#) on page **293**
- [IBM Spectrum Accelerate RPO Compliance Report](#) on page **297**
- [IBM Spectrum Virtualize RPO Compliance Report](#) on page **299**
- [NetApp ONTAP Protection Usage Report](#) on page **303**
- [NetApp ONTAP RPO Compliance Report](#) on page **306**
- [Pure Storage FlashArray RPO Compliance Report](#) on page **310**
- [Unprotected Virtual Machines Report](#) on page **313**
- [VMware RPO Compliance Report](#) on page **315**

[Storage Protection Reports](#) on page **320**

- [NetApp ONTAP OSSV Relationship Status Report](#) on page **322**
- [NetApp ONTAP SnapManager Protection Status Report](#) on page **324**
- [NetApp ONTAP Overprotected Volumes Report](#) on page **326**
- [NetApp ONTAP Qtree Protection Status Report](#) on page **328**
- [NetApp ONTAP Underprotected Volumes Report](#) on page **330**
- [NetApp ONTAP Volume Protection Status Report](#) on page **332**
- [NetApp ONTAP Transition Dependency Report](#) on page **334**

#### [Storage Utilization Reports](#) on page **336**

- [DellEMC Unity File Systems Report](#) on page **339**
- [DellEMC Unity LUNs Report](#) on page **341**
- [DellEMC Unity Pools Report](#) on page **344**
- [IBM Spectrum Accelerate Pools Report](#) on page **346**
- [IBM Spectrum Accelerate Volumes Report](#) on page **348**
- [IBM Spectrum Virtualize Consistency Groups Report](#) on page **350**
- [IBM Spectrum Virtualize Pools Report](#) on page **352**
- [IBM Spectrum Virtualize Volumes Report](#) on page **355**
- [NetApp ONTAP Aggregates Report](#) on page **358**
- [NetApp ONTAP LUNs Report](#) on page **361**
- [NetApp ONTAP Orphaned LUNs Report](#) on page **363**
- [NetApp ONTAP Quotas Report](#) on page **365**
- [NetApp ONTAP Snapshots Report](#) on page **367**
- [NetApp ONTAP Volumes Report](#) on page **369**
- [Pure Storage FlashArray Volumes Report](#) on page **371**
- [Storage Capacity Report](#) on page **374**
- [VMware Datastores Report](#) on page **375**
- [VMware LUNs Report](#) on page **377**
- [VMware Orphaned Datastores Report](#) on page **379**
- [VMware Orphaned LUNs Report](#) on page **381**
- [VMware VM Snapshot Sprawl Report](#) on page **383**
- [VMware VM Sprawl Report](#) on page **384**

## 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 catalog policy, 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 policy to summarize information about cataloged providers and the data and other resources that reside on them, then schedule the Report policy 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 **257**
- [Create a Customized Report](#) on page **259**
- [Edit a Customized Report](#) on page **260**
- [Download a Report](#) on page **261**
- [Create a Report Policy](#) on page **228**
- [Application Reports](#) on page **263**
- [System Management Reports](#) on page **270**
- [File Analytics Reports](#) on page **281**
- [Protection Compliance Reports](#) on page **291**
- [Storage Protection Reports](#) on page **320**
- [Storage Utilization Reports](#) on page **336**






## 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 policy. See [Register a Provider](#) on page 50 and [Create a NetApp ONTAP Storage Catalog Data Policy](#) on page 135.
- Run a job immediately; sometimes referred to as on demand. See [Start, Stop, and Hold a Job Session](#) on page 235.
- You can also create a schedule that can be applied to a policy so that the policy is run as defined by the parameters of the schedule. See [Create a Schedule](#) on page 119.

### 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 policy to schedule the report to run automatically. See [Create a Report Policy](#) on page 228.
- Save a report with customized parameters. See [Create a Customized Report](#) on page 259.

### RELATED TOPICS:

- [Create a Customized Report](#) on page 259
- [Create a Report Policy](#) on page 228

- [Download a Report](#) on page **261**
- [Report Overview](#) on page **256**

## 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 policy. See [Register a Provider](#) on page 50 and [Create a NetApp ONTAP Storage Catalog Data Policy](#) on page 135.

### 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 257.

### RELATED TOPICS:

- [Edit a Customized Report](#) on page 260
- [Run a Report](#) on page 257
- [Download a Report](#) on page 261
- [Report Overview](#) on page 256

## 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 policy. See [Register a Provider](#) on page **50**.
- Run a job immediately; sometimes referred to as on demand. See [Start, Stop, and Hold a Job Session](#) on page **235**.
- Create a customized report. See [Create a Customized Report](#) on page **259**.

### 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 **257**.

### RELATED TOPICS:

- [Create a Customized Report](#) on page **259**
- [Run a Report](#) on page **257**
- [Download a Report](#) on page **261**
- [Report Overview](#) on page **256**

## 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 257.

### 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.

### NEXT STEPS:




- Review your downloaded report for further analysis of your data.

### RELATED TOPICS:




- [Run a Report](#) on page 257
- [Create a Customized Report](#) on page 259

## 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 **261**
- [Run a Report](#) on page **257**

## 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 policy.

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 **264**
- [Application RPO Compliance Report](#) on page **266**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**
- [Download a Report](#) on page **261**

# 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 an Application Catalog policy. You can select one or more application provider in a single policy for cataloging. See [Create an Application Catalog Policy](#) on page 125.
- 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 96.

## Parameters

Use the following parameters to customize your report:

- Application Type: All, Microsoft Active Directory, Microsoft Exchange, Microsoft SQL, Microsoft SharePoint, Oracle. Multiple selections are supported.

The default report parameter, **All**, reports on all available application types in your configuration.

## Detail View - Microsoft

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 application server.

### Database/Mailbox Database

The name of the database or databases associated with the application server.

### Data Disk Mount Points (Volume@Storage Array)

The name of the associated database disk mount points, along with the associated volume and storage array in parentheses.

### Log Disk Mount Points

The name of the associated log disk mount points, along with the associated volume and storage array in parentheses.

### Hosting SQL Server's Instance (Microsoft SharePoint only)

The name and location of the SharePoint server's SQL server instance.



## ***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 Mount Points**

The name of the associated database disk mount points, along with the associated volume and storage array in parentheses.

### **Log Disk Mount Points**

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.

#### **RELATED TOPICS:**

- [Application Reports](#) on page **263**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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 an Application Catalog policy. You can select one or more application provider in a single policy for cataloging. See [Create an Application Catalog Policy](#) on page 125.
- Create and run an Application Copy policy. See [Create an Application Copy Policy](#) on page 146.

Use the Application RPO Compliance report to answer questions such as:

- Which of my Application Copy policies have never run successfully?

### Parameters

Use the following parameters to customize your report:

- 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.
- Display Resources 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 of the server's associated databases.

**Application Server (Type)**

The name of the application server, along with the application server type in parentheses.

**Oracle Home / SQL Instance**

The name of the Oracle Home or SQL Instance

**Policy Name**

The Copy Data policy associated with the application server.

**Last Successful Protection Time**

The most recent instance of a successful run of the Copy Data policy.

**Reason**

The reason the application server does not meet your RPO compliance parameters. Examples include no successful runs of a protection policy, 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 of the server's associated databases.

**Application Server (Type)**

The name of the application server, along with the application server type in parentheses.

**Oracle Home / SQL Instance**

The name of the Oracle Home or SQL Instance

**Database**

The names of the server's associated databases.

**Last Successful Protection Time**

The most recent instance of a successful run of the Copy policy.

**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 of the server's associated databases.

**Application Server (Type)**

The name of the application server, along with the application server type in parentheses.

**Oracle Home / SQL Instance**

The name of the Oracle Home or SQL Instance

**Source**

The source of the secondary protection.

**Destination**

The destination of the secondary protection.

**Policy Name**

The Copy Data policy associated with the application server.

**Protection Time**

The most recent instance of a successful run of the Copy policy.

**Storage Vendor**

The storage vendor associated with the application server.

**Reason**

The reason the application server does not meet your RPO compliance parameters. Examples include no successful runs of a protection policy, 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 of the server's associated databases.

**Application Server (Type)**

The name of the application server, along with the application server type in parentheses.

**Oracle Home / SQL Instance**

The name of the Oracle Home or SQL Instance

**Source**

The source of the secondary protection.

**Destination**

The destination of the secondary protection.

**Policy Name**

The Copy policy associated with the application server.

**Protection Time**

The most recent instance of a successful run of the Copy policy.

**Storage Vendor**

The storage vendor associated with the application server.

**Compliance Time Remaining**

The time remaining before your application server will be non-compliant.

**RELATED TOPICS:**

- [Application Reports](#) on page **263**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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, policies, 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 policies are associated with a Catalog?
- What is the average runtime of a policy, and when was it last successfully run?

Choose the System Management report that fits your needs:

- **Policy Report** - Review policies and their status, including the last time they ran successfully.
- **Configuration Report** - Review IBM Spectrum Copy Data Management node configuration and associated policies.
- **Catalog Summary Report** - Displays a summary of all the high level objects cataloged.
- **System Sizing Report** - View the amount of space on your IBM Spectrum Copy Data Management appliance required to catalog your NetApp ONTAP files.

### RELATED TOPICS:

- [Catalog Summary Report](#) on page **271**
- [Configuration Report](#) on page **273**
- [Policy Report](#) on page **277**
- [System Sizing Report](#) on page **279**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**
- [Download a Report](#) on page **261**

## 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, and IBM Objects**

The following fields and corresponding data display in the NetApp ONTAP Objects, VMware Objects, Application Objects, and IBM Objects, and DellEMC Unity 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 **270**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**



## 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 **270**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## Policy Report

Review the available policies in your IBM Spectrum Copy Data Management configuration. Run the Policy report to view policies by type, their average runtime, and their successful run percentage.

Use the Policy report to answer questions such as:

- What types of policies are available?
- What is the average runtime of a policy, and when was it last successfully run?
- How many times has a specific policy run successfully or failed?

### **Parameters**

Use the following parameters to edit your report:

- Policy Type  
Multiple selections are supported.
- Days Since Successful Run

The default report parameters report on all available policies.

### **Quick View**

The Quick View section displays a pie chart of the number of times a policy type successfully completed, failed, or was marked with any of the following statuses: unknown, waiting, running, stopped, partial, skipped, aborted, or stopping. Use the Policy Type parameter to display File, NetApp ONTAP, Report, Script, VMware, or all policy types.

**Note:** The Quick View section is only modified through the Policy Type parameter.

### **Summary View**

The following fields and corresponding data display in the Summary View section of the Policies report:

#### **Policy Type**

The type of policy. For example, a catalog or report policy.

#### **Policies**

The number of policies associated with the policy type.

#### **Runs**

The number of times a policy of this type ran.

#### **Completed**

The number of times a policy of this type successfully completed.

#### **Failed**

The number of times a policy of this type failed.

**Other**

The number of times a policy 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 Policies report:

**Policy**

The policy name.

**Type**

The type of policy. For example, a catalog or report policy.

**Runs**

The number of times the policy ran.

**Completed**

The number of times the policy successfully completed.

**Failed**

The number of times the policy failed.

**Other**

The number of times the policy 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 policy last ran successfully.

**Average Runtime (Days hh:mm:ss)**

The average time it takes to run the policy, listed in days, hours, minutes, and seconds.

**Success %**

The percentage of times the policy successfully completed.

**RELATED TOPICS:**

- [System Management Reports](#) on page 270
- [Report Overview](#) on page 256
- [Run a Report](#) on page 257
- [Create a Customized Report](#) on page 259

# System Sizing Report

Review the System Sizing report to view the amount of space on your IBM Spectrum Copy Data Management appliance required to catalog your NetApp ONTAP files. The report displays a list of all volumes on a NetApp ONTAP node, the number of files and snapshots on the volume, and the amount of space required to catalog the files.

## ***Parameters***

Use the following parameters to customize your report:

- NetApp ONTAP Storage
- Doc Size (KB)
- Change Rate of Files between Snapshots (%)
- Show Volume Level Detail

## ***Detail View***

The following fields and corresponding data display in the Detail View section of the System Sizing report. View additional details by enabling the Show Volume Level Detail parameter.

### **Node**

The physical server where your files are stored.

### **Volume**

The volumes on your node. The Volume field displays if the Show Volume Level Detail parameter is enabled.

### **Location**

The node where the volume resides. The Location field displays if the Show Volume Level Detail parameter is enabled.

### **No. of Files**

The number of files on the node. If the Show Volume Level Detail parameter is enabled, the number of files per volume also displays.

### **No. of Snapshots**

The number of snapshots on the node. If the Show Volume Level Detail parameter is enabled, the number of snapshots per volume also displays.

### **Space Required on Appliance (Without Snapshots)**

The space on your IBM Spectrum Copy Data Management appliance required to catalog your NetApp ONTAP files.

### **Space Required on Appliance (Including All Snapshots)**

The space on your IBM Spectrum Copy Data Management appliance required to catalog your NetApp ONTAP files based on the rate of change between snapshot copies.

**RELATED TOPICS:**

- [System Management Reports](#) on page **270**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**



## 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 policy.

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 283
- [Files By Age Report](#) on page 285

- [Files By Category Report](#) on page **287**
- [Files By Size Report](#) on page **289**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**
- [Download a Report](#) on page **261**

## 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 Catalog Data policy with the Traversal Mode set to Filewalk in conjunction with the IBM Spectrum Copy Data Management Filewalker tool. See [Create a NetApp ONTAP File Catalog Data Policy](#) on page 137.

### 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 **281**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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 Catalog Data policy with the Traversal Mode set to Filewalk in conjunction with the IBM Spectrum Copy Data Management Filewalker tool. See [Create a NetApp ONTAP File Catalog Data Policy](#) on page 137.

### **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 **281**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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 **281**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**



## 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 Catalog Data policy with the Traversal Mode set to Filewalk in conjunction with the IBM Spectrum Copy Data Management Filewalker tool. See [Create a NetApp ONTAP File Catalog Data Policy](#) on page 137.

### 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 Catalog Data policy 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 Catalog.

**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 **281**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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 policy. 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 policies 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.
- **IBM Spectrum Virtualize RPO Compliance Report** - Displays IBM storage systems in relation to your recovery point objective parameters
- **NetApp ONTAP RPO Compliance Report** - Displays the primary snapshot protection for NetApp ONTAP storage systems.
- **NetApp ONTAP Protection Usage Report** - Displays the storage usage of IBM Spectrum Copy Data Management protection policies on the volumes of your NetApp ONTAP storage systems.
- **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 **293**
- [IBM Spectrum Virtualize RPO Compliance Report](#) on page **299**
- [NetApp ONTAP Protection Usage Report](#) on page **303**
- [NetApp ONTAP RPO Compliance Report](#) on page **306**
- [Unprotected Virtual Machines Report](#) on page **313**
- [VMware RPO Compliance Report](#) on page **315**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**

- [Create a Customized Report](#) on page **259**
- [Download a Report](#) on page **261**

## 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 Storage Catalog Data policy. You can select one or more DellEMC Unity provider in a single policy for cataloging. See [Create a DellEMC Unity Catalog Data Policy](#) on page **127**.
- Create and run a DellEMC Unity Copy Data policy. See [Create a DellEMC Unity Copy Data Policy](#) on page **154**.

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 Copy Data policies 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.

#### **Policy Name**

The Copy Data policy associated with the DellEMC Unity storage system.

#### **Last Successful Protection Time**

The most recent instance of a successful run of the Copy Data policy.

#### **Reason**

The reason the DellEMC Unity storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection policy, 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.

#### **Policy Name**

The Copy Data policy associated with the DellEMC Unity storage system.

#### **Last Successful Protection Time**

The most recent instance of a successful run of the Copy Data policy.

#### **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.

#### **Policy Name**

The Replication Copy Data policy associated with the DellEMC Unity storage system.

#### **Protection Time**

The most recent instance of a successful run of the protection policy.

#### **Reason**

The reason the DellEMC Unity storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection policy, 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.

#### **Policy Name**

The Replication Copy Data policy associated with the DellEMC Unity storage system.

**Protection Time**

The most recent instance of a successful run of the Copy Data policy.

**Remaining Compliance Time**

The time remaining before your DellEMC Unity storage system will be non-compliant.

**RELATED TOPICS:**

- [Protection Compliance Reports](#) on page **291**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**



# 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 Catalog Data policy. You can select one or more IBM provider in a single policy for cataloging. See [Create an IBM Spectrum Accelerate Catalog Data Policy](#) on page **129**.
- Create and run an IBM Spectrum Accelerate Copy Data policy. See [Create an IBM Spectrum Accelerate Copy Data Policy](#) on page **157**.

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 Copy Data policies 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.

**Policy Name**

The Copy Data policy associated with the IBM storage system.

**Last Successful Protection Time**

The most recent instance of a successful run of the Copy Data policy.

**Reason**

The reason the IBM storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection policy, 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.

**Policy Name**

The Copy Data policy associated with the IBM storage system.

**Last Successful Protection Time**

The most recent instance of a successful run of the Copy Data policy.

**Compliance Time Remaining**

The time remaining before your IBM storage system will be non-compliant.

**RELATED TOPICS:**

- [Protection Compliance Reports](#) on page **291**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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 Catalog Data policy. You can select one or more IBM provider in a single policy for cataloging. See [Create an IBM Spectrum Virtualize Catalog Data Policy](#) on page **131**.
- Create and run an IBM Copy Data policy. See [Create an IBM Spectrum Virtualize Copy Data Policy](#) on page **160**.

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 Copy Data policies 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.

### **Policy Name**

The Copy Data policy associated with the IBM storage system.

### **Last Successful Protection Time**

The most recent instance of a successful run of the Copy Data policy.

### **Reason**

The reason the IBM storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection policy, 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.

### **Policy Name**

The Copy Data policy associated with the IBM storage system.

### **Last Successful Protection Time**

The most recent instance of a successful run of the Copy Data policy.

### **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.

### **Policy Name**

The Global Mirror policy associated with the IBM storage system.

### **Protection Time**

The most recent instance of a successful run of the protection policy.

### **Reason**

The reason the IBM storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection policy, 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.

### **Policy Name**

The Copy Data policy associated with the IBM storage system.

### **Protection Time**

The most recent instance of a successful run of the Copy Data policy.

### 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 **291**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## NetApp ONTAP Protection Usage Report

The NetApp ONTAP Protection Space Usage report displays the storage usage of IBM Spectrum Copy Data Management protection policies. 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 Catalog Data policy. You can select one or more NetApp ONTAP cluster providers in a single policy for cataloging. See [Create a NetApp ONTAP Storage Catalog Data Policy](#) on page 135.
- Create and run a NetApp ONTAP Copy Data policy. See [Create a NetApp ONTAP Copy Data Policy](#) on page 164.

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.
- Policy Type  
Set the IBM Spectrum Copy Data Management protection policies 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.

**Policy Names**

The names of the associated IBM Spectrum Copy Data Management protection policies.

**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.

**Policy Names**

The names of the associated IBM Spectrum Copy Data Management protection policies.

**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 291
- [Report Overview](#) on page 256



- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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 Catalog Data policy. You can select one or more NetApp ONTAP cluster providers in a single policy for cataloging. See [Create a NetApp ONTAP Storage Catalog Data Policy](#) on page 135.
- Create and run a NetApp ONTAP Copy Data policy. See [Create a NetApp ONTAP Copy Data Policy](#) on page 164.

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 Copy Data policies 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.

#### **Policy Name**

The Copy Data policy associated with the NetApp ONTAP storage system.

#### **Last Successful Protection Time**

The most recent instance of a successful run of the Copy Data policy.

#### **Reason**

The reason the NetApp ONTAP storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection policy, 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.

#### **Policy Name**

The Copy Data policy associated with the NetApp ONTAP storage system.

#### **Last Successful Protection Time**

The most recent instance of a successful run of the Copy Data policy.

#### **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.

**Policy Name**

The SnapVault or SnapMirror Copy Data policy associated with the NetApp ONTAP storage system.

**Protection Time**

The most recent instance of a successful run of the protection policy.

**Reason**

The reason the NetApp ONTAP storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection policy, 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.

**Policy Name**

The Copy Data policy associated with the NetApp ONTAP storage system.

**Protection Time**

The most recent instance of a successful run of the Copy Data policy.

**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 **291**
- [Report Overview](#) on page **256**

- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

# 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 Catalog Data policy. You can select one or more Pure Storage provider in a single policy for cataloging. See [Create a Pure Storage FlashArray Catalog Data Policy](#) on page 141.
- Create and run a Pure Storage Copy Data policy. See [Create a Pure Storage FlashArray Copy Data Policy](#) on page 167.

## 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.

**Policy Name**

The Copy Data policy associated with the Pure Storage FlashArray system.

**Last Successful Protection Time**

The most recent instance of a successful run of the Copy Data policy.

**Reason**

The reason the Pure Storage FlashArray system does not meet your RPO compliance parameters. Examples include no successful runs of a protection policy, 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.

**Policy Name**

The Copy Data policy associated with the Pure Storage FlashArray system.

**Last Successful Protection Time**

The most recent instance of a successful run of the Copy Data policy.

**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.

**Policy Name**

The Copy Data policy associated with the Pure Storage FlashArray storage system.

**Last Successful Protection Time**

The most recent instance of a successful run of the protection policy.

**Reason**

The reason the Pure Storage FlashArray storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection policy, 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.

**Policy Name**

The Copy Data policy associated with the Pure Storage FlashArray storage system.

**Last Successful Protection Time**

The most recent instance of a successful run of the Copy Data policy.

**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 **291**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**



## Unprotected Virtual Machines Report

The Unprotected Virtual Machines report displays virtual machines that are not protected with a primary snapshot or VM Copy. Quickly view all of your unprotected virtual machines, along with their eligibility for protection.

### BEFORE YOU BEGIN:

- Create and run a VMware Catalog Data policy. See [Create a VMware Catalog Data Policy](#) on page 143.
- 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 Catalog Data policy to catalog associated NetApp ONTAP storage providers. See [Create a NetApp ONTAP Storage Catalog Data Policy](#) on page 135.
- Create and run a VMware Copy Data policy. See [Create a VMware Copy Data Policy](#) on page 170.

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 **291**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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 Catalog Data policy. See [Create a VMware Catalog Data Policy](#) on page **143**.
- 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 Catalog Data policy to catalog associated NetApp ONTAP storage providers. See [Create a NetApp ONTAP Storage Catalog Data Policy](#) on page **135**.
- Create and run a VMware Copy Data policy. See [Create a VMware Copy Data Policy](#) on page **170**.

Use the VMware RPO Compliance report to answer questions such as:

- Which of my VMware Copy Data policies 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.
- Resource  
Set the resource type. 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.
- **Application Server**  
If the View VMs with Application(s) only parameter is set to Yes, set the application server types to display in the report. Available types include Active Directory Server, Exchange Server, SQL Server, and SharePoint Server. Multiple selections are supported.

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.

### **Policy Name**

The Copy Data policy 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 Copy Data policy.

### **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 policy, 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.

#### **Policy Name**

The Copy Data policy 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 Copy Data policy.

#### **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.

#### **Policy Name**

The Copy Data policy associated with the VMware object.

#### **Protection Time**

The most recent instance of a successful run of the Copy Data policy.

**Storage Vendor**

The storage vendor associated with the VMware object.

**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 secondary protection. Examples include no successful runs of a protection policy, 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.

**Policy Name**

The Copy Data policy associated with the VMware object.

**Protection Time**

The most recent instance of a successful run of the Copy Data policy.

**Storage Vendor**

The storage vendor associated with the VMware object.

**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.

**RELATED TOPICS:**

- [Protection Compliance Reports](#) on page **291**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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 policy. 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 SnapManager Protection** - Review NetApp ONTAP SnapManager protection status for SQL and Exchange applications.
- **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 **322**
- [NetApp ONTAP SnapManager Protection Status Report](#) on page **324**
- [NetApp ONTAP Overprotected Volumes Report](#) on page **326**
- [NetApp ONTAP Qtree Protection Status Report](#) on page **328**
- [NetApp ONTAP Underprotected Volumes Report](#) on page **330**
- [NetApp ONTAP Volume Protection Status Report](#) on page **332**
- [NetApp ONTAP Transition Dependency Report](#) on page **334**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**
- [Download a Report](#) on page **261**

## 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 320
- [Report Overview](#) on page 256

- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## NetApp ONTAP SnapManager Protection Status Report

Run the NetApp ONTAP SnapManager Protection Status Report to review the SnapManager backup status of SQL and Exchange applications. Use the Secondary Protection parameter to review SnapVault or SnapMirror recovery point details for the volume associated with the application server.

Use the NetApp ONTAP SnapManager Protection Status report to answer questions such as:

- How many recovery points are associated with a specific application server?
- What is the latest snapshot creation time of a specific application server?

### **Parameters**

Use the following parameter to customize your report:

- Application Type  
Multiple selections are supported.
- Application Server  
Multiple selections are supported.
- Secondary Protection Type

Secondary recovery point names display in the report if secondary SnapVault or SnapMirror protection for the volume is available. The default report parameters report on all application types and servers with secondary SnapVault protection if available.

### **Recovery Points**

The following fields and corresponding data display in the Recovery Points section of the NetApp ONTAP SnapManager Protection Status report:

#### **Application Server**

The name of the application server.

#### **Application Type**

The application type. For example, SQL or Exchange.

#### **Node**

The NetApp ONTAP storage location associated with the application server.

#### **Volume**

The volume associated with the application server.

#### **Location**

The location where the volume resides on the application sever.

#### **Recovery Points**

The number of recovery points on the volume.

**Latest Recovery Point**

The name of the most recent recovery point.

**Latest Snap Creation Time**

The date and time of the most recent SnapManager backup.

**SnapMirror / SnapVault**

The name of the secondary recovery point if secondary SnapVault or SnapMirror protection for the volume is available. This field is controlled by the Secondary Protection parameter.

**RELATED TOPICS:**

- [Storage Protection Reports](#) on page **320**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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.

#### **Host**

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 **320**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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.

#### Node

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 **320**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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.

#### **Node**

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 **320**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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.

#### Node

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.

**Node**

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 **320**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

# 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 Catalog Data policy to catalog your NetApp ONTAP storage systems operating in 7-mode. You can select one or more NetApp ONTAP cluster provider in a single policy for cataloging. See [Create a NetApp ONTAP Storage Catalog Data Policy](#) on page **135**.
- Create and run a NetApp ONTAP Copy Data policy. See [Create a NetApp ONTAP Copy Data Policy](#) on page **164**.

## 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:

### Node

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:

**Node: 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 **320**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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 policy.

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.
- **IBM Spectrum Accelerate Pools** - 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** - View the total space used by and free space available to your IBM Spectrum Accelerate volumes.
- **IBM Spectrum Virtualize Consistency Groups** - 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.
- **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.

### 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.



- **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** - 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 the storage capacity of your IBM Spectrum Virtualize pools, DellEMC Unity pools, and NetApp ONTAP aggregates.
- **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.

## 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 **339**
- [DellEMC Unity LUNs Report](#) on page **341**
- [DellEMC Unity Pools Report](#) on page **344**
- [IBM Spectrum Accelerate Pools Report](#) on page **346**
- [IBM Spectrum Accelerate Volumes Report](#) on page **348**
- [IBM Spectrum Virtualize Consistency Groups Report](#) on page **350**
- [IBM Spectrum Virtualize Pools Report](#) on page **352**
- [IBM Spectrum Virtualize Volumes Report](#) on page **355**
- [NetApp ONTAP Aggregates Report](#) on page **358**
- [NetApp ONTAP LUNs Report](#) on page **361**
- [NetApp ONTAP Orphaned LUNs Report](#) on page **363**
- [NetApp ONTAP Quotas Report](#) on page **365**
- [NetApp ONTAP Snapshots Report](#) on page **367**
- [NetApp ONTAP Volumes Report](#) on page **369**
- [Pure Storage FlashArray Volumes Report](#) on page **371**
- [Storage Capacity Report](#) on page **374**
- [VMware Datastores Report](#) on page **375**
- [VMware LUNs Report](#) on page **377**
- [VMware Orphaned Datastores Report](#) on page **379**
- [VMware Orphaned LUNs Report](#) on page **381**
- [VMware VM Snapshot Sprawl Report](#) on page **383**
- [VMware VM Sprawl Report](#) on page **384**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**
- [Download a Report](#) on page **261**

## 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 Catalog Data policy. You can select one or more DellEMC Unity providers in a single policy for cataloging. See [Create a DellEMC Unity Catalog Data Policy](#) on page 127.

### 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.

#### 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 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.

**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 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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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 Catalog Data policy. You can select one or more DellEMC Unity providers in a single policy for cataloging. See [Create a DellEMC Unity Catalog Data Policy](#) on page 127.

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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## DellEMC Unity Pools Report

Run the DellEMC Unity Pools report to display 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 DellEMC Unity pools.

### BEFORE YOU BEGIN:

- Create and run a DellEMC Unity Catalog Data policy. You can select one or more DellEMC Unity providers in a single policy for cataloging. See [Create a DellEMC Unity Catalog Data Policy](#) on page 127.

Use the DellEMC 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 DellEMC 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 DellEMC Unity Host parameter.

### Summary View

The following fields and corresponding data display in the Summary View section of the DellEMC Unity Pools report:

#### DellEMC Unity Storage

The physical server where your files are stored.

#### Storage Pools

The number of storage pools on the DellEMC 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 336
- [Report Overview](#) on page 256
- [Run a Report](#) on page 257
- [Create a Customized Report](#) on page 259

# 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 Catalog Data policy. You can select one or more IBM providers in a single policy for cataloging. See [Create an IBM Spectrum Accelerate Catalog Data Policy](#) on page 129.

## 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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

# 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 Catalog Data policy. You can select one or more IBM provider in a single policy for cataloging. See [Create an IBM Spectrum Accelerate Catalog Data Policy](#) on page 129.

## 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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

# 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 Catalog Data policy. You can select one or more IBM providers in a single policy for cataloging. See [Create an IBM Spectrum Virtualize Catalog Data Policy](#) on page 131 and [Create an IBM Spectrum Protect Snapshot Catalog Data Policy](#) on page 133.

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.

### IBM Storage

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.

**IBM Storage**

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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

# 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 Catalog Data policy. You can select one or more IBM providers in a single policy for cataloging. See [Create an IBM Spectrum Virtualize Catalog Data Policy](#) on page 131.

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:

### IBM Storage

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.

**IBM Storage**

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 MDisk 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 MDisk 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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

# 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 Catalog Data policy. You can select one or more IBM provider in a single policy for cataloging. See [Create an IBM Spectrum Virtualize Catalog Data Policy](#) on page 131.

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
- Record Limit
- Show Flash Copies

## 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:

### IBM Storage

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.

**Allocated**

The portion of real capacity that is being used to store data.

**Available**

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.

**IBM Storage**

The physical storage system where your files are stored.

**Storage Pool**

The volume's associated storage pool.

**Capacity**

The volume storage capacity that is available to a host.

**Available**

The difference between the real capacity and used capacity values for volume copies.

**Allocated**

The portion of real capacity that is being used to store data.

**Status**

The status of the volume. A volume can be online, offline, or degraded.

**Warning**

For thin provisioned or compressed volume copies, a warning is generated at this percentage of the volume capacity.

**Thin Provisioned**

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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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.

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:

#### **Node**

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.

**Node**

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 **336**
- [Report Overview](#) on page **256**

- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**



## 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.

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:

#### **Node**

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.

### **Node**

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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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.

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:

#### **Node**

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.

### **Node**

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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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.

### ***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:

#### **Node**

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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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.

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:

#### **Node**

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.

### **Node**

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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**



## 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.

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:

#### **Node**

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.

**Node**

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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

# 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 Catalog Data policy. You can select one or more Pure Storage providers in a single policy for cataloging. See [Create a Pure Storage FlashArray Catalog Data Policy](#) on page 141.

## ***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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

# 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 Catalog Data policies. See [Register a Provider](#) on page 50 and [Plan Overview](#) on page 99.

## Parameters

Use the following parameter to customize your report:

- Storage Vendor

## Detail View

The following fields and corresponding data display in the Detail 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.

## RELATED TOPICS:

- [Storage Utilization Reports](#) on page 336
- [Report Overview](#) on page 256
- [Run a Report](#) on page 257
- [Create a Customized Report](#) on page 259

## 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.

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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**



## 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.

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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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.

### **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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## VMware Orphaned LUNs Report

Review VMware orphaned LUNs. These are the LUNs that are neither used as datastores nor RDMS.

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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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.

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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

## 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.

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 **336**
- [Report Overview](#) on page **256**
- [Run a Report](#) on page **257**
- [Create a Customized Report](#) on page **259**

# Maintenance

[Maintenance Overview](#) on page **388**

[Log On to the Virtual Appliance](#) on page **389**

[Set Time Zone](#) on page **390**

[Collect Logs For Troubleshooting](#) on page **391**

[Manage the Virtual Appliance](#) on page **393**

[Update IBM Spectrum Copy Data Management](#) on page **394**

[Install the Marketplace RPM](#) on page **396**

[Upload an SSL Certificate](#) on page **397**

## Maintenance Overview

In most cases, IBM Spectrum Copy Data Management is installed on a virtual appliance. The virtual appliance contains the application and the Catalog.

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 setting the time zone, collecting logs, updating the application, and reviewing the configuration of the virtual appliance.

### RELATED TOPICS:

- [Log On to the Virtual Appliance](#) on page **389**
- [Set Time Zone](#) on page **390**
- [Collect Logs For Troubleshooting](#) on page **391**
- [Manage the Virtual Appliance](#) on page **393**
- [Upload an SSL Certificate](#) on page **397**
- [LDAP User Name Syntax](#) on page **408**

## 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 36.

### 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 391
- [Manage the Virtual Appliance](#) on page 393
- [Set Time Zone](#) on page 390

## Set Time Zone

Access the IBM Spectrum Copy Data Management virtual appliance through vSphere Client to set your local time zone. The time zone change functionality is built into the appliance. The time zone is used for policy scheduling.

### 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 **36**.

### To set a time zone:

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 **Set Timezone**.
4. Select a time zone. Scroll through the list by pressing Shift+Page Up and Shift+Page Down.

### RELATED TOPICS:

- [Log On to the Virtual Appliance](#) on page **389**
- [Manage the Virtual Appliance](#) on page **393**

## 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.

### 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](http://www.hypexr.org/linux_scp_help.php).

### 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/ECX/tools/scripts**.  
Run **logcollect** using the **sudo** command. For example, at the command prompt enter:  
**\$ sudo ./logcollect**  
or  
**\$ sudo /opt/ECX/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 **Monitor**  tab. For example, at the command prompt enter:  
**\$ sudo ./logcollect -job <job ID>**  
or  
**\$ sudo /opt/ECX/tools/scripts/logcollect -job <job ID>**
4. Find the collected logs archived as a zip file in **/opt/ECX/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](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 389
- [Monitor a Job Session](#) on page 237



## Manage the Virtual Appliance

Log on to the web-based management 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 36.

### To manage the virtual machine:

1. From a supported browser, enter the following URL:  
*https://<HOSTNAME>:5480/*  
where *<HOSTNAME>* is the IP address of the virtual machine where the application is deployed.
2. In the login window, enter your user name and password to access the management console of the virtual machine. The default user name is **administrator** and 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 36

## 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 36.
- 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.

### To update your IBM Spectrum Copy Data Management appliance with an upgrade package:

1. From a machine with Internet Access, download the zipped update file from IBM Fix Central at <https://www-945.ibm.com/support/fixcentral>
2. Log into the command prompt at IBM Spectrum Copy Data Management with username **administrator** and the password **ecxadLG235**.
3. Copy the zipped update file to the appliance's **/tmp** directory.
4. Execute “cd /tmp” to move to the **tmp** directory.
5. Execute “mkdir update” to create a directory named **update**.
6. Execute “cd update” to move to the update directory.
7. If the update is delivered as a .zip file, unzip the file using the following command: “unzip ../<update-file-name.zip>”.  
**Note:** Execute “yum install unzip” to install Unzip if not already present in the installation.
8. From a supported web browser, access the IBM Spectrum Copy Data Management web-based management console at the following address: <https://<IP Address>:5480>. The username is **administrator** and the password is **ecxadLG235**.
9. Navigate to the **Update** tab, click **Settings**, then click **Use Specified Repository**.
10. In the Repository URL field, enter “file:///tmp/update”.
11. Perform the following sequence under **Actions** in the upper right corner:
  1. Click **Save Settings**
  2. Click **Status**, then click **Check Updates**
  3. Once the update version displays, click **Install Updates**
12. After the update completes, navigate to the **System** tab, then click **Reboot** under Actions in the upper right corner.

**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 **36**
- [Manage the Virtual Appliance](#) on page **393**

## 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 36.

### To set up 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 navigate to "/" through the **cd** command, then execute one of the following commands:
  - "rpm -ivh marketplaceadmin-<version>.noarch.rpm" to install the RPM package if it is not currently installed
  - "rpm -Uvh marketplaceadmin-<version>.noarch.rpm" to update an installed RPM package
  - "rpm -qa" to list all packages installed on the system
  - "rpm -q marketplaceadmin-<version>.noarch.rpm" to verify if the package is installed
  - "rpm --verify marketplaceadmin-<version>.noarch.rpm" to verify that the installed package is proper

### 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 36
- [Manage the Virtual Appliance](#) on page 393

## 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 36.
- 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 408.

### 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>:5480/*  
where *<HOSTNAME>* is the IP address of the virtual machine where the application is deployed.
4. In the login window, enter your user name and password to access the management console of the virtual machine. The default user name is **administrator** and the default password is **ecxadLG235**.
5. Click the **Certificate** tab.
6. In the Certificate file field, browse for the certificate file on your computer, then click **Import**. A successful upload message appears.  
**Note:** You may need to enter your user name and password after clicking **Import**.
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 **50**
- [LDAP User Name Syntax](#) on page **408**
- [User Administration and Security Management](#) on page **14**

# Documentation and Support


[Documentation Roadmap](#) on page **400**


[About the Help System](#) on page **401**

# 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.

By default, the Help system accessed from the management interface is online. To ensure you are seeing the latest Help system, clear your browser cache before selecting Help. To use a local Help system that does not require an Internet connection click the arrow next to **Help**  to open the About IBM Spectrum Copy Data Management window. Select **Use local help system**. The local version might not contain the latest updates.

## *Quick Start Guide*

This PDF is intended for users and administrators. It contains a product overview, a summary of the workflow, and the steps to deploy and start the application, register and catalog resources, run a job, perform searches, and generate reports. Additionally, it contains checklists for application deployment.

This PDF is also intended for the Super User and system administrators for IBM Spectrum Copy Data Management. It provides procedures on how to install, set up, and start the application, and lists the system requirements pertinent to successfully install and effectively run the application. It also provides procedures to add and maintain Users including configuring LDAP authentication.

## *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, and schedule jobs.



### **RELATED TOPICS:**


- [About the Help System](#) on page **401**



## 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.

By default, the Help system accessed from the management interface is online. To ensure you are seeing the latest Help system, clear your browser cache before selecting Help. To use a local Help system that does not require an Internet connection click the arrow next to **Help**  to open the About IBM Spectrum Copy Data Management window. Select **Use local help system**. The local version might not contain the latest updates.

**Note:** 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 policy returns results for schedule and policy as well as "schedule policy."
- Enter words separated with AND, +, or & to find topics that contain all of the words separated by the operators. For example, schedule AND policy, schedule + policy, or schedule & policy.
- Enter words separated with OR to find topic that contain any of the words separated by OR. For example, schedule OR policy.

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 **400**

# Reference Topics

[Search and Filter Guidelines](#) on page **404**

[Select, Sort, and Reorder Columns](#) on page **407**

[LDAP User Name Syntax](#) on page **408**

[Return Code Reference](#) on page **409**

## 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 **241**
- [View Object Details](#) on page **246**
- [View NetApp ONTAP File Details](#) on page **247**
- [Create a NetApp ONTAP File Catalog Data Policy](#) on page **137**

## 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 **63**
- [Edit a Schedule](#) on page **121**
- [Edit a Policy](#) on page **230**
- [Monitor a Job Session](#) on page **237**
- [Search Overview](#) on page **240**
- [Run a Report](#) on page **257**

## 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 **397**
- [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 Policy](#) on page 225

# 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 **412**.

For questions related to resources used by IBM Spectrum Copy Data Management, see [Resources](#) on page **413**.

For questions related to connecting to IBM Spectrum Copy Data Management resources, see [Connectivity](#) on page **414**.

For questions related to IBM Spectrum Copy Data Management cataloging and jobs, see [Cataloging](#) on page **415**.

For questions related to general IBM Spectrum Copy Data Management operation, see [Operation](#) on page **416**.

For questions related to Copy Data and Use Data policies in IBM Spectrum Copy Data Management, see [Copy Data/Use Data](#) on page **417**.

## ***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 **36**.

### **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.

### **Is it important to configure the time zone for the IBM Spectrum Copy Data Management virtual appliance?**

Yes. The time zone is for the scheduling of policies. An incorrect time zone may cause scheduled tasks to run at a shifted time. See the topic [Set Time Zone](#) on page 390.

### **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 20 and [Install IBM Spectrum Copy Data Management as a Virtual Appliance](#) on page 36.

**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 20 and [Start IBM Spectrum Copy Data Management](#) on page 38.

**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 6.8 is the operating system on the IBM Spectrum Copy Data Management virtual appliance.

**Is Java required for accessing IBM Spectrum Copy Data Management?**

A browser that supports JavaScript is required.

**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 Copy Data and Use Data support for customers with VMware leveraging heterogeneous storage, extending Copy Data 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 a catalog policy, does that automatically discover all the ESX servers within that vCenter?**

Yes. Once cataloged, view available VMware resources through the catalog 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 Monitor tab?**

The number of operations in progress on the Monitor 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 Monitor 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 job runs for a given policy, older objects for that job are purged from the Catalog. This retention parameter is set when the policy is defined.

The Maintenance policy removes resources and associated objects created by IBM Spectrum Copy Data Management when a policy 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 Policy](#) on page **232**.

## *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 **249**.

The recommended Best Practice for protection/recovery is to use Copy Data and Use Data policies 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 do I access the most current Help?**

By default, the Help system accessed from the management interface is online. To ensure you are seeing the latest Help system, clear your browser cache before selecting Help.

To use a local Help system that does not require an Internet connection click the arrow next to **Help**  to open the About IBM Spectrum Copy Data Management window. Select **Use local help system**. The local version might not contain the latest updates.

### **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 391.

### **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, click **Download Audit Log**, then select a location to save the log file.

## ***Copy Data/Use Data***

### **For a Copy Data policy, how do I update the retention after a policy has run?**

Open the existing policy, click Snapshot in the workflow pane, and update the Keep Snapshots parameter. The retention policy changes to the supplied value when the policy is next run.

### **To what extent do the IBM Spectrum Copy Data Management Copy Data and Use Data functions impact the performance of NetApp storage systems?**

The Copy Data and Use Data 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 Copy Data policy, can I exclude swap partitions?**

No. The lowest granularity of protection is a virtual machine.

# Acronyms

## A

---

### AD

Active Directory

### API

Application Programming Interface

## 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

## 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 Catalog, accessible through VMware vSphere client. The appliance is also referred to as the virtual appliance or virtual machine.

## C

---

### **Catalog**

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 Catalog resides on the appliance.

### **Catalog Browser**

A feature enabling you to explore cataloged resources and their underlying objects and object metadata. Drill into the Catalog 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.

### **Catalog Data Policy**

A policy for gathering and recording objects and object metadata about specified resources.

### **clone mode**

A Use Data Policy 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 are located in your IT environment, and leverage the most appropriate copy for any given use case.

### **Copy Data Policy**

A policy that leverages Copy Data Management technology for replicating and intelligently reusing snapshots, vaults, and mirrors.

## 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 policies.

**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.

---

**J**

---

**job**

An execution of a policy.

---

**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.

**monitor**

A feature enabling you to monitor job sessions and execute certain job control functions. From the monitor, you can view all scheduled and unscheduled job sessions, and start, stop, and hold jobs.

---

## N

---

**NetApp provider**

A NetApp storage system such as a FAS.

---

## O

---

**object**

A resource that has metadata about it stored in the Catalog. Examples of objects are files, directories, qtrees, and volumes.

---

## P

---

**policy**

A defined set of tasks and rules that govern the execution of a job. A policy can be applied to one or more jobs. IBM Spectrum Copy Data Management includes support for the following policy types: Catalog Data policies, Copy Data policies, Use Data policies, Analyze Data policies, Script policies.

**production mode**

A Use Data Policy 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 by using a catalog data policy.

**report policy**

A policy that runs predefined or customized reports summarizing information about objects and object metadata managed by cataloged resources. A report policy is a type of Analyze Data policy, 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.

**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 policy**

A policy that runs a set of commands on the IBM Spectrum Copy Data Management appliance. The script policy 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 Catalog 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 Catalog 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.

**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 Use Data Policy mode where temporary virtual machines are created in a fenced network for such uses as



development/testing, snapshot verification, and disaster recovery verification.

## U

---

### **Use Data Policy**

A policy that leverages Copy Data Management technology for reusing and recovering resources from snapshots, vaults, and mirrors.

### **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 copy data and use data policies. A workflow is defined visually by the storage administrator in the management interface.

# Trademarks

## ***Commonly Used Company and Product Names***

Companies and products listed here may be used in the documentation:

### **Adobe®**

PDF

### **DellEMC®**

EMC®, EMC<sup>2</sup>®, 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 401  
Acronyms 419  
Add a Key 93  
Add a Resource Pool 74  
Add a Role 77  
Add a Site 44  
Add Credentials to a Virtual Machine 69  
Appliance Protection 416  
Application Configuration Report 264  
Application RPO Compliance Report 266  
Application Workflow 11  
Audience and Purpose 9  
Authentication 14  
AWS Requirements 35

## B

---

Best Practice 99, 119, 121, 137, 180, 215, 261-262  
Best Practices for Configuring Tenants 89  
Browse Catalog 252  
Browser Support 20

## C

---

Catalog Browser 240, 252  
Catalog Policy 99

Catalog Protection 416  
Catalog Summary Report 271  
Character Strings 404  
Clone 218  
Collect Logs For Troubleshooting 391  
Columns 407  
Configuration Report 273  
Configure  
    Node 50  
    Overview 48  
    Provider 50  
Configure Accounts 79  
Configure Resource Pools 74  
Configure Roles 77  
Configure Storage Workflows 102  
Configure Tenants 87  
Copy Data 417  
Create 175  
    Customized Report 259  
    NetApp Copy Data Policy 164  
    NetApp File Catalog Data Policy 137  
    NetApp Storage Catalog Data Policy 135  
    Report Policy 228  
    Schedule 119  
    Script Policy 225  
    VMware Copy Data Policy 170  
    VMware Use Data Policy 213  
Create a NetApp Use Data Policy 201  
Create a Pure Storage FlashArray Copy Data Policy 167

Create a VMware Catalog Data Policy 143  
Create an Application Catalog Policy 125  
Create an Application Copy Policy 146  
Create an Application Use Policy 179  
Create an IBM XIV Copy Data Policy 157  
Credentials 69  
Customized Report 259-260

## D

---

Dashboard Overview 40  
Definitions 421  
Delete  
    Catalog Policy 231  
    Generated Report 262  
    Schedule 123  
Delete a Resource Pool 75  
Delete a Role 78  
Delete an Account 80  
Deploy the Application 36  
Deployment Checklist 19  
Documentation Roadmap 400  
Documentation Suite 400  
Download  
    Report 261  
    Search Results 251

## E

---

Edit  
    Customized Report 260

    Policy 230  
    Schedule 121  
Edit a Provider 66  
Edit a Resource Pool 75  
Edit a Role 78  
Edit an Account 80  
EMC Unity 48, 50  
EMC VNX Catalog Data Policy Requirements 22  
EMC VNX Copy and Use Data Policy Requirements 26  
Encryption 15  
End IV 218  
Export Search Results 251

## F

---

FAQ 412  
FCM 22  
Fibre Channel 181, 183, 217, 219  
File Analytics Reports 281  
File Search 241  
Files By Age Report 285  
Files By Category Report 287  
Files By Size Report 289  
Files Usage by Owner Report 283  
Filter Guidelines 404  
Find and Restore a File 249  
Frequently Asked Questions 412

---

**G**

---

Generate a Report 257  
Generated Report 261  
Glossary 421

---

**H**

---

Help 400-401  
Host 48

---

**I**

---

IBM Storage and IBM FCM Catalog Data Policy  
Requirements 22, 26  
Identification and Authentication 14  
iGroup 202, 215  
Install the Application 36  
Installation Requirements 20  
Internationalization 24  
iSCSI 27, 202, 215

---

**J**

---

Job 99  
Job History 237  
Job Monitor 237  
Job Session 235, 237

---

**K**

---

Knowledge Base articles  
1010992 30

1013246 30

---

**L**

---

LDAP  
Authentication 14, 19, 408  
Resources 63  
Server 48, 50-51, 63, 415  
User Name Syntax 408  
Legal 426  
Load Balancing 175  
Load Sharing 175  
Log On to the Virtual Appliance 389  
Logs For Troubleshooting 391

---

**M**

---

Maintenance Overview 388  
Maintenance Policy 232  
Manage the Virtual Appliance 393  
Manuals 400  
Marketplace 396  
Monitor a Job Session 237  
Monitor Overview 234

---

**N**

---

Native User 14  
NetApp 7-Mode 48, 50  
NetApp Cluster Mode 48, 50  
NetApp Copy Data and Use Data Policy  
Requirements 26

NetApp File Catalog Data Policy Requirements 23  
NetApp Knowledge Base articles  
    1010992 30  
    1013246 30  
NetApp Protection Usage Report 303  
NetApp RPO Compliance Report 306  
NetApp Storage Catalog Data Policy Requirements 23  
NetApp Storage System 48, 50  
NetApp Volumes 63  
Node 48

---

## O

Object Search 241  
Online Help 400-401  
Oracle 59  
Oracle Requirements 31  
OSSV Relationship Status Report 322  
Overview 10

---

## P

Pattern Matching 404  
PDFs 400  
Performance Considerations 175  
Plan Overview 99  
Policy Report 277  
Policy Types 99  
Ports 15  
Predefined Report 256-257

Protection Compliance Reports 291  
Providers 48  
Provisioned User 14  
Proxy 175

---

## R

Recovery Catalog 240, 252  
Register a Provider 50  
Register a Resource 50  
Registrations 426  
Report Browser 257, 259-260  
Report Policy 100  
Reports Overview 256  
Resource Browser 51-53, 55-59, 61, 63, 66, 68  
Return Code Reference 409  
Role-Based Access Control Overview 72  
RRP 218  
Run a Report 257

---

## S

sAMAccountName 51  
Schedule 238  
Script Policy 100, 225  
Search for Objects 241  
Search Guidelines 404  
Search Overview 240  
Security 14  
Select, Sort, and Reorder Columns 407  
Set Time Zone 390

SMTP Host 48, 50  
Start, Stop, and Hold a Job Session 235  
Storage Protection Reports 320  
Storage System 48, 50  
Storage Utilization Reports 336  
Subnet 216  
System Management Reports 270  
System Requirements 20  
System Sizing Report 279

## T

---

Tabs 64, 244, 253  
Terminology 421  
Time Machine 240, 252  
Time Zone 390, 413  
Trademarks 426  
Trigger 119  
Troubleshooting Logs 391

## U

---

Unprotected Virtual Machines Report 313  
Unregister a Provider 68  
Upload an LDAP SSL Certificate 397  
Use Data 417  
User Administration and Security Management 14  
User Data Security 14  
User Provisioning 48, 50  
Using State and Status Arguments in  
Postscripts 222

## V

---

VADP 175  
VADP Proxy 175  
VASA 215  
vCenter Host 48, 50  
View a Provider 63  
View File Details 247  
View Object Details 246  
Virtual Appliance 36, 389, 393  
Virtual Machine Privileges 27  
Virtual Machines 63  
VMFS 170, 214  
VMware Catalog 240, 252  
VMware Catalog Data Policy Requirements 24  
VMware Copy Data and Use Data Policy  
Requirements 27  
VMware Datastores Report 375  
VMware Hosts 63  
VMware LUNs Report 377  
VMware Orphaned Datastores Report 379  
VMware Orphaned LUNs Report 381  
VMware RPO Compliance Report 315  
VMware Server 48, 50, 54  
VMware VM Snapshot Sprawl Report 383  
VMware VM Sprawl Report 384  
VVOL 102, 171, 215



## **W**

---

Wildcard Guidelines 404

Workflow 11