

IBM Storage Copy Data Management
2.2.24

Installation and User's Guide



Note:

Before you use this information and the product it supports, read the information in [“Notices” on page 345.](#)

This edition applies to version 2, release 2, modification 24 of IBM® Storage Copy Data Management (product numbers 5737-B34, 5641-CD4, 5641-CD5, 5641-CD6) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2017, 2024.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

- About this publication.....ix**
 - Audience and purpose.....ix

- What's new in IBM Storage Copy Data Management 2.2.24..... xi**

- Chapter 1. IBM Storage Copy Data Management overview..... 1**

- Chapter 2. Getting off to a quick start.....3**
 - Start IBM Storage Copy Data Management.....3
 - Register a provider.....4
 - Create an inventory job definition.....4
 - Run a job.....4
 - Search for objects.....4
 - Generate reports.....5
 - Backup.....5
 - Restore.....5

- Chapter 3. User administration and security management..... 7**

- Chapter 4. Installation and setup..... 11**
 - Deployment Checklist.....11
 - Access and default credentials..... 12
 - VMware vSphere Privileges 14
 - Installing IBM Storage Copy Data Management as a virtual appliance..... 16
 - Installing the IBM Storage Sentinel Security Scan..... 18
 - Uninstalling the agent from Windows-based application server..... 18
 - Starting IBM Storage Copy Data Management..... 19

- Chapter 5. Dashboard..... 21**

- Chapter 6. Sites and providers..... 23**
 - Sites and providers Overview..... 23
 - Adding a site.....24
 - Editing a site.....24
 - Deleting a site..... 25
 - Registering a provider..... 25
 - Registering an Application Server - File System..... 26
 - Registering an Application Server - InterSystems Database..... 27
 - Registering an Application Server - Oracle.....28
 - Registering an Application Server - SAP HANA..... 29
 - Registering an Application Server - SQL..... 30
 - Registering an IBM Storage Virtualize provider.....31
 - Registering an LDAP provider.....32
 - Registering a NetApp ONTAP provider..... 33
 - Registering multiple NetApp ONTAP providers through the Discover feature..... 34
 - Registering a Pure Storage FlashArray provider.....35
 - Registering a Security Scan Server in IBM Storage Copy Data Management..... 36
 - Registering an SMTP provider..... 37
 - Registering a VMware provider..... 38

Registering multiple VMware providers through the Discover feature.....	39
Viewing a provider.....	40
Viewing a list of providers.....	41
Viewing job definitions associated with providers.....	42
Editing a provider.....	43
Unregistering a provider.....	44
Adding credentials to a virtual machine.....	45
Adding credentials for a single virtual machine.....	45
Adding credentials for multiple virtual machines.....	45
Chapter 7. Requiring multifactor authentication.....	47
Multifactor authentication for the ADMIN user.....	47
Creating multifactor authentication user account.....	48
Enabling time-based one-time multifactor authentication for a user.....	48
Setting up a multifactor authentication user account in the browser.....	48
Expiring time-based one-time secret key of a user account.....	49
Disabling multifactor authentication.....	49
Chapter 8. Role-based access control.....	51
Role-based access control overview.....	51
Configure resource pools.....	52
Configure roles.....	54
Configure accounts.....	55
VMware admin role-based access control configuration.....	57
Database admin role-based access control configuration.....	58
NetApp ONTAP admin role-based access control configuration.....	59
IBM admin role-based access control configuration.....	60
Pure Storage FlashArray admin role-based access control configuration.....	61
Chapter 9. Configure tenants.....	63
Adding a tenant.....	63
Editing a tenant.....	64
Deleting a tenant.....	64
Chapter 10. Best practices for configuring tenants.....	65
Chapter 11. Identities.....	67
Identities overview.....	67
Adding a key.....	67
Adding an SSH key through the Generate a keypair for me method and registering an associated provider.....	67
Adding an SSH key through the I will provide a keypair method and registering an associated provider.....	68
Adding a credential.....	70
Editing a credential.....	70
Deleting a credential.....	70
Chapter 12. Configure SLA policies.....	73
Creating an SLA policy.....	73
Configure IBM Storage Virtualize SLA policies.....	74
Configure IBM Storage Virtualize for Snapshot SLA policies.....	78
Configure NetApp ONTAP SLA policies.....	79
Configure Pure Storage FlashArray SLA policies.....	82
Chapter 13. Configure scripts.....	85
Uploading a script.....	85

Replacing a script.....	85
Deleting a script.....	86
Chapter 14. Schedules.....	87
Creating a schedule.....	87
Editing a schedule.....	88
Deleting a schedule.....	89
Chapter 15. Jobs.....	91
Jobs overview.....	91
Start, pause, and hold a job session.....	91
Starting a job session.....	91
Pausing and resuming a job session.....	92
Holding and releasing a job session.....	92
Canceling a job session.....	93
Monitor a job session.....	93
Monitoring a running job session.....	93
Filtering the list of jobs based on type or status.....	94
Viewing information about specific job sessions.....	94
Job definition overview.....	95
Inventory jobs.....	97
Creating an Inventory job Definition - Database.....	97
Creating an Inventory job Definition - File System.....	99
Creating an Inventory job Definition - IBM Storage Virtualize.....	100
Creating an Inventory job Definition - IBM Storage Virtualize for Snapshot.....	102
Creating an Inventory job Definition - NetApp ONTAP Storage.....	103
Creating an Inventory job Definition - NetApp ONTAP File.....	104
Creating an Inventory job Definition - Pure Storage FlashArray.....	107
Creating an Inventory job Definition - VMware.....	109
Backup jobs.....	110
Creating a backup job definition - InterSystems Database.....	110
Creating a scanning backup job definition - InterSystem IRIS-based application (on AIX).....	112
Creating a backup job definition - SAP HANA.....	114
Creating a backup job definition - Oracle.....	116
Creating a scanning backup job definition - Oracle-based application (on AIX).....	121
Creating a backup job definition - SQL.....	125
Creating a backup job definition - File System.....	129
Creating a backup job definition - IBM Storage Virtualize.....	131
Creating a backup job definition - IBM Storage Virtualize for Snapshot.....	136
Creating a backup job definition - NetApp ONTAP.....	140
Creating a backup job definition - Pure Storage FlashArray.....	142
Creating a backup job definition - VMware.....	145
Creating VMware Backup job proxies.....	149
Restore jobs.....	151
Creating a Restore job definition - InterSystems Database.....	151
Creating a Restore job definition - SAP HANA.....	154
Creating a Restore job definition - Oracle.....	157
Creating a Restore job definition - SQL.....	162
Creating a Restore job definition - File System.....	166
Creating a Restore job definition - IBM Storage Virtualize.....	169
Creating a Restore job definition - IBM Storage Virtualize for Snapshot.....	176
Creating a Restore job definition - NetAPP ONTAP.....	179
Creating a Restore job definition - Pure Storage FlashArray.....	187
Creating a Restore job definition - VMware.....	191
Restore Jobs - Rename mount points and initialization parameter options.....	205
Using state and status arguments in postscripts.....	206
System jobs.....	207

Maintenance job.....	207
Creating a report job definition.....	208
Creating a script job definition.....	209
Editing a job definition.....	211
Deleting a job definition.....	212
Chapter 16. Searching.....	213
Searching overview.....	213
Searching for objects.....	214
Viewing object details.....	218
Viewing NetApp ONTAP file details.....	219
Finding and restoring a file.....	220
Downloading search results.....	221
Browsing inventory.....	222
Browsing cataloged providers.....	222
Browsing through the Inventory on a previous date.....	223
Chapter 17. Report.....	225
Report Overview.....	225
Running a report.....	226
Creating a customized report.....	227
Editing a customized report.....	228
Downloading a Report.....	229
Deleting a generated report.....	230
Application reports.....	230
Application configuration report.....	231
Application RPO compliance report.....	233
System management reports.....	236
Catalog summary report.....	237
Configuration report.....	238
Job report.....	240
Job sessions report.....	242
File analytics reports.....	243
File usage by owner report.....	244
Files by age report.....	246
Files by category report.....	247
Files by size report.....	249
Protection compliance reports.....	250
File System RPO compliance report.....	252
IBM Storage Virtualize RPO compliance report.....	254
NetApp ONTAP protection usage report.....	257
NetApp ONTAP RPO compliance report.....	259
Pure Storage FlashArray RPO compliance report.....	261
Recovery points report.....	264
Unprotected Virtual Machines report.....	265
VMware RPO compliance report.....	266
Storage protection reports.....	269
NetApp ONTAP OSSV relationship status report.....	271
NetApp ONTAP overprotected volumes report.....	272
NetApp ONTAP Qtree protection status report.....	273
NetApp ONTAP underprotected volumes report.....	274
NetApp ONTAP volume protection status report.....	275
NetApp ONTAP transition dependency report.....	277
Storage utilization reports.....	278
IBM Storage Virtualize Consistency Groups Report.....	282
IBM Storage Virtualize Pools Report.....	283
IBM Storage Virtualize Volumes Report.....	285

Instant Disk Restore Volumes Report.....	287
NetApp ONTAP Aggregates Report.....	288
NetApp ONTAP LUNs Report.....	290
NetApp ONTAP Orphaned LUNs Report.....	292
NetApp ONTAP Quotas Report.....	293
NetApp ONTAP Snapshots Report.....	295
NetApp ONTAP Volumes Report.....	296
Pure Storage FlashArray Volumes Report.....	298
Storage Capacity Report.....	300
VM and Storage Mapping Report.....	301
VMware Datastores Report.....	302
VMware LUNs Report.....	304
VMware Orphaned Datastores Report.....	306
VMware Orphaned LUNs Report.....	307
VMware VM Snapshot Sprawl Report.....	309
VMware VM Sprawl Report.....	310
VMware VM Storage Report.....	312
Chapter 18. Maintenance.....	315
Maintenance overview.....	315
Log on to the virtual appliance.....	316
Setting the time zone.....	316
Collecting logs for troubleshooting.....	317
Collecting audit logs from the support menu.....	317
Collecting logs from the Support menu.....	318
Collecting the IBM Storage Copy Data Management logs from the virtual appliance.....	319
Modifying job log options.....	320
Updating global settings.....	321
Managing the administrative console.....	322
Pre-upgrade resource adjustments.....	322
Update IBM Storage Copy Data Management.....	324
Updating IBM Storage Copy Data Management appliance from 2.2.18 or later versions to 2.2.23 or later.....	324
Updating IBM Storage Copy Data Management appliance from 2.2.16 or 2.2.17 version to 2.2.18 or later versions.....	325
Updating IBM Storage Copy Data Management appliance from 2.2.15 or earlier versions to 2.2.16 or later versions.....	326
Backup and Restore the Catalog.....	328
Modifying the network settings.....	329
Uploading an SSL certificate.....	330
Horizontal scale out.....	331
Deploy a MongoDB virtual machine (VM) from the .ova file.....	332
Copying the appropriate VMDKs from IBM Storage Copy Data Management to the MongoDB VM	332
Configuring the MongoDB VM.....	333
Configuring IBM Storage Copy Data Management to use the MongoDB VMs rather than the collocated MongoDB processes.....	333
Binding directories to ensure proper function of log collection.....	334
Restoring a snapshot from a FlexGroup volume to another FlexGroup volume.....	334
Chapter 19. Documentation and support.....	337
Documentation roadmap.....	337
About the help system.....	337
Chapter 20. Reference topics.....	339
Search and filter guidelines.....	339
Select, sort, and reorder columns.....	341
LDAP username syntax.....	341

Return code reference.....	342
Notices.....	345
Glossary.....	349

About this publication

This publication provides overview, planning, installation, and user instructions for IBM Storage Copy Data Management.

Audience and purpose

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

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

What's new in IBM Storage Copy Data Management 2.2.24

New features and enhancements are available in IBM Storage Copy Data Management 2.2.24.

General Enhancements

Feature	Description	Benefit
Support for SAP HANA 2.0 SPS 07	SAP HANA 2.0 SPS 07 can now be backed up and restored.	Allows users to continue protecting their databases when upgrading to the latest version of SAP HANA.
Support to scan VMware virtual machines for ransomware	IBM Storage Copy Data Management can now scan the VMware virtual machines for ransomware.	Expands ransomware detection support to the VMware virtual machines to improve cyber-resilience.
Support to perform thin-clone and normal restore for IBM Storage Virtualize for Snapshot	The user can now configure the restore type for IBM Storage Virtualize for Snapshot to be either Normal or Thin-clone by using a property file.	Added flexibility to configure the nature of the mounted volumes from snapshot to be either Thin-clone or Normal.
Platform update including security patches	The underlying software components that are used in IBM Storage Copy Data Management is upgraded to use the latest version.	Software vulnerabilities in the appliance have been patched.

Chapter 1. IBM Storage Copy Data Management overview

IBM Storage Copy Data Management is a Copy Data Management (CDM) platform that can bring modernization to an existing environment without disruption.

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

IBM Storage Copy Data Management delivers “in-place” copy data management to enterprise storage arrays from IBM, NetApp and Pure Storage, allowing the IT team to use its existing infrastructure and data in a manner that is efficient, automated, scalable, and easy to use. IBM Storage Copy Data Management modernizes IT processes, enables key use cases, and does it all without additional hardware.

For the latest system requirements for IBM Storage Copy Data Management, see [Requirements](#).

Chapter 2. Getting off to a quick start

IBM Storage Copy Data Management workflow includes, starting IBM Storage Copy Data Management, registering a provider, cataloging data, searching for objects, generating reports, and copying and using data.

Related concepts

[“Sites and providers overview” on page 23](#)

A site is a user-defined grouping of providers that is generally based on location to help quickly identify and interact with data that is created through Copy Data Management jobs. Sites are assigned when registering providers. When creating Backup and Restore jobs, sites clearly identify where your data is replicated by location. Providers are physical servers that host objects and attributes. Once a provider is registered in IBM Storage Copy Data Management, cataloging, searching, and reporting can be performed.

[“Start, pause, and hold a job session” on page 91](#)

From the Jobs tab you can run a job session on demand, pause or cancel a running job, and hold all future scheduled instances of a job from running until you are ready for the job to proceed.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Start IBM Storage Copy Data Management” on page 3](#)

Start IBM Storage Copy Data Management to begin using the application and its features.

Related reference

[“Searching for objects” on page 214](#)

Use the search feature to find objects on providers that are cataloged in IBM Storage Copy Data Management. Examples of objects are volumes, files, snapshots, qtrees, and virtual machines. You can tailor your search by applying filters.

Start IBM Storage Copy Data Management

Start IBM Storage Copy Data Management to begin using the application and its features.

Procedure

To Start IBM Storage Copy Data Management, complete the following steps:

1. In 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 Storage Copy Data Management.

2. Enter your username and password, which is provided by the IBM Storage Copy Data Management System Administrator.

If this is your first time logging on to IBM Storage Copy Data Management as a Super User, the default username is `admin` and the default password is `password`. You are prompted to reset the default Super User password.

3. Click **Sign In**. The application starts.

Register a provider

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

Procedure

To register a provider, complete the following steps:

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

Create an inventory job definition

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

Procedure

To create an inventory job definition, complete the following steps:

1. Click the **Jobs** tab and the **policies** view.
2. Select a provider type, click **New**, and then select **Inventory**.
3. Select one or more resources to catalog from the list of available providers.
4. Select the options for your job. Also, enter notifications. If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
5. Optionally, select one or more defined schedules for your job and save the job definition.

Run a job

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

Procedure

To run a job, complete the following steps:

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

Search for objects

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

Procedure

To search for objects, complete the following steps:

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.

Procedure

To generate reports, complete the following steps:

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

Backup

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

Procedure

To backup your data, complete the following steps:

1. Click the **Jobs** tab.
2. Select a provider type, click **New** and then select **Backup**.
3. Select providers to copy or protect, as well as a storage workflow that meets your copy data criteria.
4. Complete the job definition including notification and other options. Save the job definition.
5. Run the backup job. The selected source data is copied to the destination in accordance with the defined job parameters.

Restore

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

Procedure

To restore, complete the following steps:

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

Chapter 3. User administration and security management

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

IBM Storage Copy Data Management security objectives are:

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

Remember: IBM Storage 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 username and password. System Administrators have the option of adding native users or importing groups of provisioned users through LDAP authentication. Native usernames are not case-sensitive. LDAP username case sensitivity relies on the configuration of your LDAP server.

User Data Security

IBM Storage 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 Storage Copy Data Management functionality.

Sensitive data is encrypted when stored.

Data in transit is also protected. IBM Storage Copy Data Management protects the confidentiality of the user and system credentials. Sensitive data is encrypted or transported by using TLS and HTTPS. The user login is protected via HTTPS for browser client to IBM Storage 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 Storage Copy Data Management identifies the following types of sensitive data: native user credentials, IBM, NetApp, and Pure Storage FlashArray storage system credentials, VMware/ESX host credentials, and user credentials.

Security Management

Security management identifies the interfaces that manage the security functions in the IBM Storage 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 Storage Copy Data Management to use LDAP, and configuring IBM Storage Copy Data Management to use HTTPS.

Following are the security management functions in IBM Storage Copy Data Management:

- Adding, editing, and deleting a user
- Configuring authentication mode
- Assigning roles to a user
- 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 timeout period that is assigned that is for the application in minutes. If the user does not refresh or request a window within the timeout period, the session ends automatically. Session timeout is set for 30 minutes and cannot be changed.
- Users are uniquely identified by entering a username and password.
- Role-based access control is employed. Once a user is added to IBM Storage 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 Storage 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 by using TLS and HTTPS. User credentials such as passwords are safely stored in the IBM Storage Copy Data Management database. Obtaining and storing this sensitive data constitutes the basic function of the IBM Storage Copy Data Management application. This data is subject to the user data security requirements.

Ports

Port	Service	Comment
22	OpenSSH 5.3 (protocol 2.0)	Port open within the firewall
25	smtp, non-TLS connection for Simple Mail Transfer Protocol	Service used by IBM Storage Copy Data Management
68	bootpc in DHCP clients, DHCP Listener UDP	Service used by IBM Storage Copy Data Management
80/443	http/https	Service used by IBM Storage Copy Data Management
389	LDAP, non-TLS connection for Lightweight Directory Access Protocol	Service used by IBM Storage Copy Data Management
443	smtp, TLS connection for Simple Mail Transfer Protocol	Service used by IBM Storage Copy Data Management
636	LDAP, TLS connection for Lightweight Directory Access Protocol	Service used by IBM Storage Copy Data Management
1433	sql, SQL Service	Service used by IBM Storage Copy Data Management

Port	Service	Comment
4369	epmd, Erlang port mapper	Service used by IBM Storage Copy Data Management
5432	postgresql, PostgreSQL DB 8.4.1-8.4.4	Service used by IBM Storage Copy Data Management
5480	ssl/http, vami	Port open within the firewall, IBM Storage Copy Data Management 2.2.5 and earlier
5985	smtp, TLS connection for Simple Mail Transfer Protocol	WinRM, Windows Remote Management
8090	admin console, IBM Storage Copy Data Management Administrative Console	Port open within the firewall
8092	adminconsole, IBM Storage Copy Data Management Administrative Console	Port open within the firewall, IBM Storage Copy Data Management 2.2.6 only
8443	ssl/http, Apache Tomcat/Coyote JSP engine 1.1	Port open within the firewall
8761	Discovery Server	Service used by . Locates registered microIBM Storage Copy Data Management services.
9090	Liberty Server	Service used to serve the Knowledge Center documentation
27017	MongoDB mongod	Service used by IBM Storage Copy Data Management
27018	MongoDB mongod	Service used by IBM Storage Copy Data Management
55672	rabbitMQ, RabbitMQ administrative	Service used by IBM Storage Copy Data Management

Related concepts

“Registering a provider” on page 25

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

“Role-based access control overview” on page 51

Role-based access control allows you to set the resources and permissions available to IBM Storage Copy Data Management accounts.

Chapter 4. Installation and setup

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

Deployment checklist

This checklist is for IBM Storage Copy Data Management deployment to a VMware appliance host.

Following are the pre-deployment, deployment, and post-deployment procedures.

Step	Action	Related Topic	✓
1	Reference the locations of access and default credentials for IBM Storage Copy Data Management.	Access and default credentials	

Step	Action	Related Topic	
1	Install IBM Storage Copy Data Management by deploying an OVF template to create a virtual appliance containing the application on a VMware host.	Installing IBM Storage Copy Data Management as a virtual appliance	
2	To access IBM Storage Copy Data Management, upload a valid product key to the virtual machine where IBM Storage Copy Data Management is deployed.	Installing IBM Storage Copy Data Management as a virtual appliance	
3	Launch IBM Storage Copy Data Management to set a new Super User password.	Start IBM Storage Copy Data Management	
4	Configure LDAP Authentication and establish a secure connection to the LDAP server.	Registering a provider	
5	Create accounts through role-based access control.	Role-based access control overview	

Step	Action	Related Topic	
1	Launch IBM Storage Copy Data Management to begin using the application and its features..	Start IBM Storage Copy Data Management	
2	Update IBM Storage Copy Data Management to keep it current with new features, enhancements, and upgrades.	Update IBM Storage Copy Data Management	
3	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.	Managing the administrative console	

Access and default credentials

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

In general, when you log in to IBM Storage Copy Data Management for the first time for each of these accounts, you will be prompted to change the default password. New passwords must adhere to the following rules:

- Be a minimum of 15 characters in length
- Contain at least one uppercase letter
- Contain at least one lowercase letter
- Contain at least one digit
- Contain at least one special character: (!, @, #, \$, %, etc.)
- Three or more characters in the new password that must not be present in the old password.

Important: When you deploy IBM Storage Copy Data Management, you must change the **SYSTEM** account password through SSH before you can access the Administrative Console. SSH into the appliance with the username `administrator` and the default password. After authenticating, you will be prompted to change the password from the default. This account can now be used to log into the Administrative Console.

Table 4.

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

Table 4. (continued)

Interface	Location	Default Credentials	Notes®
IBM Storage Copy Data Management Administrative Console	https:// <HOSTNAME_OR_IP>:8090/	Username: administrator Password: ecxadLG235	This interface is used to make administrative changes to the IBM Storage Copy Data Management server. In the drop down menu, select System to use the administrator account. Otherwise, select IBM Storage Copy Data Management to enter a valid IBM Storage Copy Data Management account username and password.
IBM Storage Copy Data Management Command Line / Console	SSH or Console	Username: administrator Password: ecxadLG235	The administrator account may be used to access the IBM Storage Copy Data Management server through SSH or through console access. This account may also be used to access the Administrative Console (above). Note: This account does not have sudo access. If elevated privileges are required, use the root account.
IBM Storage Copy Data Management Command Line / Console	SSH or Console	Username: root Password: ecxDP758	The root account may be used to make changes to the IBM Storage Copy Data Management server through SSH or through console access. Use this account when elevated privileges are required.

Table 4. (continued)

Interface	Location	Default Credentials	Notes®
IBM Storage Copy Data Management Agent User	Application Server - SSH or Console	Suggested Username:: ecxagent	When adding a database or file system, you must create an agent user on that system. For more information, see the appropriate topic, "Sample Configuration of an IBM Storage Copy Data Management Agent User" for the file system or application. This account is only needed on the client systems, not on the IBM Storage Copy Data Management OVA.

VMware vSphere Privileges

In IBM Storage Copy Data Management, the user account that is associated with the provider is not assigned the Administrator role for an inventory object. Instead, the user must be assigned to a role that has, at a minimum, the following privileges. These privileges are propagated to child objects. For more information about adding a permission to an inventory object, refer to the VMware documentation.

In the following list, the bold text indicates the vCenter Server Object, and the indented text with a bullet is the require privilege of that object. Some entries listed with the version to indicate different levels.

vCenter Server Object	Required Privileges
Alarm	<ul style="list-style-type: none"> • Acknowledge alarm • Set alarm status
Cryptographic Operations (6.5 and 6.7)	<ul style="list-style-type: none"> • Add disk • Direct access • Encrypt • Encrypt new • Manage encryption policies
Datastore	<ul style="list-style-type: none"> • Allocate space • Browse datastore • Low-level file operations • Remove datastore • Remove file • Update virtual machine files

vCenter Server Object	Required Privileges
Distributed switch	<ul style="list-style-type: none"> • Port configuration operation • Port setting operation
Folder	<ul style="list-style-type: none"> • Create folder
Global	<ul style="list-style-type: none"> • Cancel task
Host > Configuration	<ul style="list-style-type: none"> • Storage partition configuration
Inventory Service > Tagging (6.0) vSphere Tagging (6.5, and 6.7)	<ul style="list-style-type: none"> • Assign or Unassign vSphere Tag • Create vSphere Tag • Create vSphere Tag Category • Modify UsedBy Field for Category • Modify UsedBy Field for Tag
Network	<ul style="list-style-type: none"> • Assign network
Resource	<ul style="list-style-type: none"> • Apply recommendation • Assign a vApp to resource pool • Assign virtual machine to resource pool • Migrate powered off virtual machine • Migrate powered on virtual machine • Query vMotion
Virtual Machine > Configuration	<ul style="list-style-type: none"> • Add existing disk • Add new disk • Add or remove device • Advanced (6.0 and 6.5) • Advanced configuration (6.7) • Change CPU count • Change memory (6.7) • Configure raw device (6.7) • Disk change tracking (6.0 and 6.5) • Memory (6.0 and 6.5) • Modify device settings • Raw device (6.0 and 6.5) • Reload from path • Remove disk • Rename • Settings • Toggle disk change tracking (6.7)
Virtual Machine > Guest Operations	<ul style="list-style-type: none"> • Guest Operation Modifications • Guest Operation Program Execution • Guest Operation Queries

vCenter Server Object	Required Privileges
Virtual Machine > Interaction	<ul style="list-style-type: none"> • Backup operation on virtual machine • Power[®] Off • Power On
Virtual Machine > Inventory	<ul style="list-style-type: none"> • Register • Remove • Unregister
Virtual Machine > Provisioning	<ul style="list-style-type: none"> • Allow read-only disk access • Mark as template • Mark as virtual machine
Virtual Machine > Snapshot management	<ul style="list-style-type: none"> • Create snapshot • Remove snapshot • Revert snapshot
vApp	<ul style="list-style-type: none"> • Add virtual machine • Assign resource pool • Assign vApp • Create • Delete • Power Off • Power On • Rename • Unregister • vApp resource configuration

Installing IBM Storage Copy Data Management as a virtual appliance

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

Before you begin

You must complete a new installation of IBM Storage Copy Data Management when migrating from a previous version of IBM Storage Copy Data Management to version 2.2.19.0. This requires the catalog to be transferred from the previous version of IBM Storage Copy Data Management to the new deployment of IBM Storage Copy Data Management 2.2.19.0. For more information, follow the procedure in [“Transferring the catalog from a previous version of IBM Storage Copy Data Management”](#) on page 326.

- 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 Storage 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 Storage Copy Data Management deploys, you must redeploy the virtual machine.
- For later versions of vSphere, the vSphere Web Client may be required to deploy IBM Storage Copy Data Management appliances.

Procedure

To deploy IBM Storage Copy Data Management as a virtual appliance, complete the following steps:

1. Use the vSphere Client to deploy IBM Storage Copy Data Management. From the **Actions** menu, choose **Deploy OVF Template...** The Deploy OVF template window opens.
2. Enter the URL for the IBM Storage Copy Data Management OVA template file or specify the location of the OVA template file and select it. Click **Next**.
3. Provide a name for the virtual machine in the **Virtual machine name** field and select a folder to contain the virtual machine. Click **Next**.
4. Identify the compute resource to be used by the virtual machine. Select the destination compute resource and then click **Next**.
5. Review the advanced configuration details.

When deploying IBM Storage Copy Data Management in a VMware vCenter 7.0 environment or later, you might receive a warning stating that the certificate is not trusted. The certificate that is supplied with IBM Storage Copy Data Management is valid. Click **Ignore** to close the warning in the vCenter OVA UI wizard when you deploy IBM Storage Copy Data Management.

If you prefer that your VMware vCenter fully validate the OVA certificate, download the chain and certificate authority (CA) certificates used for the signing of the OVA and install these in the vCenter. You must extract the chain and root certificates directly from DigiCert. The DigiCert Assured ID Root CA certificate and the DigiCert SHA2 Extended Validation Server CA certificate can be used to validate the OVA code signed image. Download the DigiCert Assured ID Root CA certificate at <https://cacerts.digicert.com/DigiCertAssuredIDRootCA.crt.pem> and the DigiCert SHA2 Extended Validation Server CA certificate at <https://cacerts.digicert.com/DigiCertSHA2ExtendedValidationServerCA.crt.pem>.

- a) Download each certificate in Privacy Enhanced Mail (PEM) format.
- b) Log in to the vCenter using the VMware administrator account.
- c) Navigate to Administration > Certificate Management.
- d) Locate Trusted Root Certificates and click the Add link to add the two downloaded certificates.

After the two certificates are added, you can import the OVA without having to click the ignore link in the UI.

Click **Next**.

6. Read and accept the End User License Agreement. Select **I accept all license agreements** if you agree and then click **Next**.
7. Select the storage location for the configuration and disk files. 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. 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**.
8. Select networks for the deployed template to use. Several available networks on the ESXi 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**.
9. Enter network properties for the virtual machine. Enter the hostname in the **Hostname** field. If no hostname is supplied, *localhost* is used. Enter network information in **Default Gateway, DNS,**

Network IP address and **Network Prefix**. 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**.

10. Review your template selections. Click **Finish** to exit the wizard and to start deployment of the OVF template. Deployment may take several minutes.
11. After the template deployment completes, power on your newly created virtual machine. The virtual machine must remain powered on for the IBM Storage Copy Data Management application to be accessible.
12. Make a note of the hostname or IP address of the newly created virtual machine. This is needed to log on to the application. Find the hostname or 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 to initialize completely.

What to do next

- Set your local time zone. See [“Setting the time zone”](#) on page 316.
- After the first use, you can enable additional users to logon by linking to an LDAP server.
- Start IBM Storage Copy Data Management and begin using it from any supported web browser. See [“Start IBM Storage Copy Data Management”](#) on page 3.

Related tasks

[“Start IBM Storage Copy Data Management”](#) on page 3

Start IBM Storage Copy Data Management to begin using the application and its features.

Installing IBM Storage Sentinel Security Scan

This topic section cover installing IBM Storage Sentinel Security Scan.

To install IBM Storage Sentinel Security Scan, do the steps that follow:

- Install the IBM Storage Sentinel Security Scan, refer to [IBM Storage Sentinel](#).
- Add a license to the IBM Storage Sentinel Security Scan, refer to [Final engine setup](#).



Attention: A federation consists of two or more IBM Storage Sentinel Security Scan software systems together to share work. Install the IBM Storage Sentinel Security Scan on the federation manager machines and then add a license to the manager. Now you can install new member IBM Storage Sentinel Security Scan software servers and connect them to the appropriate manager machines. Member machines share the license from the manager, as they do not receive a standalone license for installation. To create a federation, on IBM Storage Copy Data Management you must register federation manager first and then register the member.

Uninstalling the agent from Windows-based application server

It may be necessary to uninstall the IBM Storage Copy Data Management agent from Windows-based application servers. This procedure is required to uninstall the agent if you choose to remove the application server from protection.

Before you begin

You will need administrative access to the application server from which the agent is to be uninstalled.

Procedure

1. Using Remote Desktop Protocol (RDP), log in to the application server that is running agent that is to be uninstalled.
2. Click **Start** and type `services.msc`. The Services window opens.
3. Locate the service that contains the agent service in the **Name** column.

4. Right-click and select **Stop** to stop the agent service.
5. Click **Start** and type cmd. Select **Run as Administrator** to launch the Command Prompt with administrative privileges. Select Yes for any prompts. The Administrator: Command Prompt window opens.
6. At the prompt, navigate to the directory.

```
CD "\\Program Files\IBM\IBM Storage Copy Data Management\bin"
```

7. Remove the agent.

```
ecxagent.exe -r
```

8. Next, navigate back to the Program Files directory and clean up any remaining directory structure to remove files related to the agent.

```
CD "\\Program Files"  
RD /S "IBM Storage Copy Data Management"
```

Results

The agent service is no longer running and is uninstalled. Any supporting files for the agent are also removed.

Starting IBM Storage Copy Data Management

Start IBM Storage Copy Data Management to begin using the application and its features.

Before you begin

- IBM Storage Copy Data Management must be installed before starting the application. See [“Installing IBM Storage Copy Data Management as a virtual appliance”](#) on page 16.
- The System Administrator must provide you with the IP address for the virtual appliance and the IBM Storage Copy Data Management username and password.

Procedure

To start IBM Storage Copy Data Management, complete the following steps:

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 Storage Copy Data Management.

2. In the logon dialog, enter your username and password, which is provided by the IBM Storage Copy Data Management System Administrator.

If this is your first time logging on to IBM Storage Copy Data Management as a Super User, the default username is `admin` and the default password is `password`. You are prompted to reset the default Super User password.

3. Click **Sign In**. The application starts.

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

What to do next

- After the first use, enable additional users to log on by adding native users or linking to an LDAP server. See [“Role-based access control overview”](#) on page 51.

- Add storage systems and virtual machine resources to the IBM Storage Copy Data Management database. See [“Registering a provider”](#) on page 25 and [“Jobs overview”](#) on page 91.
- Search or browse for objects that match certain criteria. See [“Searching for objects”](#) on page 214 and [“Browsing inventory”](#) on page 222.
- Generate reports with predefined or customized parameters. See [“Report Overview”](#) on page 225.

Related tasks

[“Installing IBM Storage Copy Data Management as a virtual appliance”](#) on page 16

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

Chapter 5. Dashboard

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

Before you begin, review IBM Storage Copy Data Management features and data flow information. See [IBM Storage Copy Data Management Overview](#).

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

Available Dashboard Widgets

Inventory Statistics

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

Appliance Filesystems

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

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

What to do next

- 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 concepts

[“Sites and providers overview” on page 23](#)

A site is a user-defined grouping of providers that is generally based on location to help quickly identify and interact with data that is created through Copy Data Management jobs. Sites are assigned when registering providers. When creating Backup and Restore jobs, sites clearly identify where your data is replicated by location. Providers are physical servers that host objects and attributes. Once a provider is registered in IBM Storage Copy Data Management, cataloging, searching, and reporting can be performed.

[“Role-based access control overview” on page 51](#)

Role-based access control allows you to set the resources and permissions available to IBM Storage Copy Data Management accounts.

[“Jobs overview” on page 91](#)

You can create and edit job definitions from the **Jobs** tab. You can also start, monitor, stop, and resume job sessions from the Jobs tab. From this pane, you can also view all scheduled and unscheduled job sessions, and start jobs before scheduled run times.

[“Monitor a job session” on page 93](#)

You can view the details of a job session that is currently running or one that has finished.

[“Searching overview” on page 213](#)

Use IBM Storage 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.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Chapter 6. Sites and providers

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

Sites and providers overview

A site is a user-defined grouping of providers that is generally based on location to help quickly identify and interact with data that is created through Copy Data Management jobs. Sites are assigned when registering providers. When creating Backup and Restore jobs, sites clearly identify where your data is replicated by location. Providers are physical servers that host objects and attributes. Once a provider is registered in IBM Storage Copy Data Management, cataloging, searching, and reporting can be performed.

Supported provider types are:

- **Application servers:** Supported application database types include InterSystems Caché and InterSystems IRIS (collectively referred to as InterSystems Database), Oracle, SAP HANA, and SQL. Use the File System application type to register file systems for physical servers running Windows, Linux®, and AIX®. See associated application requirements for supported storage types.
- **IBM storage systems:** Supported types include IBM Storage 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 Storage Virtualize 7.3 and later/8.1.2 and later, including IBM SAN Volume Controller, IBM Storwize®, and IBM FlashSystem® V9000 and 9100 systems. Supported SQL Server versions include SQL Server 2012, SQL Server 2014, SQL Server 2016 standalone and AlwaysON. Supported operating system platforms include Windows 2012R2, Windows 2016 running on vSphere VM using VMDK configuration.
- **Pure Storage systems:** Supported storage platforms include Pure Storage FlashArray.
- **SMTP hosts.** Register an SMTP server to enable email notifications from IBM Storage 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 Storage Copy Data Management, such as jobs or customized reports, will have full access to interact with those providers or resources regardless of role-based access control restrictions. For example, if a user's permission allows them to register NetApp providers, they will also be able to view, edit, and unregister the NetApp providers that they registered, even if the necessary permissions are not assigned to them through role-based access control.

Related concepts

[“Registering a provider” on page 25](#)

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

[“Viewing a provider” on page 40](#)

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.

Related tasks

[“Editing a provider” on page 43](#)

Revise the properties of a provider as needed.

[“Unregistering a provider” on page 44](#)

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

Adding a site

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

Before you begin

- Click **View Relationship** to view the resources that are assigned to the site.
- Review the properties and location of your current providers. See [“Sites and providers overview” on page 23](#) for a list of supported providers.

Procedure

To add a site, complete the following steps:

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

What to do next

Assign sites to new and existing providers. See [“Registering a provider” on page 25](#) and [“Editing a provider” on page 43](#).

Related tasks

[“Editing a site” on page 24](#)

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

[“Deleting a site” on page 25](#)

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

Editing a site

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

Procedure

To edit a site, complete the following steps:

1. Click the **Configure** tab. On the Views pane, select **Sites & Providers**, then select the **Sites** tab. The Sites pane opens.
2. In the Sites pane, select the **Site** to edit by clicking in the row containing the site name.
3. Click **Edit**. The Edit Site dialog opens.
4. Revise the site name and description.

5. To set this site as the default site, select **Set as default**. New storage providers are automatically assigned to the default site unless another site is selected.
6. Click **OK**. The revisions are applied to the site.

What to do next

Assign sites to new and existing providers. See [“Registering a provider” on page 25](#) and [“Editing a provider” on page 43](#).

Related tasks

[“Adding a site” on page 24](#)

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

[“Deleting a site” on page 25](#)

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

Deleting a site

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

Procedure

To delete a site, complete the following steps:

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

What to do next

Assign sites to new and existing providers. See [“Registering a provider” on page 25](#) and [“Editing a provider” on page 43](#).

Related tasks

[“Adding a site” on page 24](#)

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

[“Editing a site” on page 24](#)

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

Registering a provider

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

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

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

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

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

Registering an Application Server - File System

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 [“Adding a site” on page 24](#).
- When registering providers it is recommended to assign all related sources, such as hosting vCenters and related storage systems, to the same Site.
- For application server it is recommended to keep all resources, such as hosting vCenters and related storage systems, configured in the same Site.

Procedure

To register an Application Server-File System, complete the following steps:

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

Site

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

Name

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

Host Address

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

Port

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

OS Type

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

System Credential

Select or create your file system or volume's credentials. See [“Identities overview” on page 67](#).

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

Notice: 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.

Registering an Application Server - InterSystems Database

You can register InterSystems Caché and InterSystems IRIS application servers in IBM Storage Copy Data Management as an InterSystems Database.

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 [“Adding a site” on page 24](#).
- When registering providers it is recommended to assign all related sources, such as hosting vCenters and related storage systems, to the same Site.
- For application server it is recommended to keep all resources, such as hosting vCenters and related storage systems, configured in the same Site.

Important: Avoid hosting multiple servers on the common data stores. Doing so may result in conflicts for restore jobs such as using SAP HANA's revert functionality.

Procedure

To register an Application Server-InterSystems Database, complete the following steps:

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

Site

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

Name

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

Host Address

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

Port

The communications port of the provider you are adding.

Run Inventory job after registration

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

Type

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

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

OS Type

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

Authentication

IBM Storage Copy Data Management connects to the InterSystems Database server as a local operating system user through an SSH key or password. See [“Identities overview” on page 67](#).

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

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

System Credential

Select or create your InterSystems Database credentials. See [“Identities overview”](#) on page 67.

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

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

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

Registering an Application Server - Oracle

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 [“Adding a site”](#) on page 24.
- When registering providers it is recommended to assign all related sources, such as hosting vCenters and related storage systems, to the same Site.
- For application server it is recommended to keep all resources, such as hosting vCenters and related storage systems, configured in the same Site.

Procedure

To register an Application Server-Oracle, complete the following steps:

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

Site

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

Name

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

Host Address

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

Run Inventory job after registration

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

Type

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

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

Authentication

IBM Storage 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 67.

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 Storage Copy Data Management first confirms a network connection and then adds the provider to the database.

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

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

Registering an Application Server - SAP HANA

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 [“Adding a site”](#) on page 24.
- When registering providers it is recommended to assign all related sources, such as hosting vCenters and related storage systems, to the same Site.
- For application server it is recommended to keep all resources, such as hosting vCenters and related storage systems, configured in the same Site.
- **Important:** To successfully catalog a SAP HANA application provider, you must also register associated vCenters.

Procedure

To register an Application Server-SAP HANA, complete the following steps:

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

Site

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

Name

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

Host Address

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

Port

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

Run Inventory job after registration

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

vCenter

The vCenter location of the SAP HANA application server.

System Credential / Database Credential(s)

Select or create your SAP HANA credentials. See [“Identities overview”](#) on page 67.

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

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

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

Registering an Application Server - SQL

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 [“Adding a site”](#) on page 24.
- When registering providers it is recommended to assign all related sources, such as hosting vCenters and related storage systems, to the same Site.
- For application server it is recommended to keep all resources, such as hosting vCenters and related storage systems, configured in the same Site.
- When registering physical SQL servers it is recommended to register via the DNS server. The IBM Storage Copy Data Management appliance must be resolvable and route-able by the DNS server; the physical SQL server will communicate back to IBM Storage Copy Data Management through DNS.
- Note that for physical SQL servers, you must allow outgoing connections to port 8443 on the IBM Storage Copy Data Management appliance from the SQL server.
- **Note:** To successfully catalog a virtual SQL application provider, you must also register associated vCenters.

Procedure

To register an Application Application Server - SQL, complete the following steps:

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

Site

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

Name

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

Host Address

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

Port

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

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

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

Run Inventory job after registration

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

Type

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

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

vCenter

The vCenter location of the virtual SQL application server.

System Credential:

Select or create your SQL credentials. See [“Identities overview” on page 67](#).

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

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

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

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

Registering an IBM Storage Virtualize provider

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 [“Adding a site” on page 24](#).
- When registering providers it is recommended to assign all related sources, such as hosting vCenters and related storage systems, to the same Site.
- For application server it is recommended to keep all resources, such as hosting vCenters and related storage systems, configured in the same Site.

Procedure

To register an IBM Storage Virtualize provider, complete the following steps:

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

4. Populate the fields in the dialog:

Site

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

Name

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

Host Address

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

Comment

Optional provider description.

Run Inventory job after registration

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

Credentials

Select or create your IBM Storage Virtualize credentials. See [“Identities overview”](#) on page 67.

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

5. Click **OK**. IBM Storage 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 Storage Copy Data Management.

Registering an LDAP provider

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 [Adding a site](#)
- When registering providers it is recommended to assign all related sources, such as hosting vCenters and related storage systems, to the same Site.
- For application server it is recommended to keep all resources, such as hosting vCenters and related storage systems, configured in the same Site.

Procedure

To register an LDAP provider, complete the following steps:

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

Name

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

Host Address

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

Port

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

Use SSL

Enable to establish a secure connection to the LDAP server.

Credentials

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

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

Base DN

The location where users and groups can be found.

User Filter

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

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

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

Note that this entry also controls the type of user name that appears in IBM Storage 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 Storage 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.

Registering a NetApp ONTAP provider

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 [“Adding a site” on page 24](#).
- When registering providers it is recommended to assign all related sources, such as hosting vCenters and related storage systems, to the same Site.
- For application server it is recommended to keep all resources, such as hosting vCenters and related storage systems, configured in the same Site.

Procedure

To register a NetApp ONTAP provider, complete the following steps:

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

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

Site

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

Name

A user-defined name for the 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 TLS** check box to enable an encrypted Secure Socket Layer connection. The typical default port is 80 for non TLS connections or 443 for TLS connections.

Comment

Optional provider description.

Run Inventory job after registration

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

Credentials

Select or create your NetApp ONTAP credentials. See [“Identities overview”](#) on page 67.

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

5. Click **OK**. IBM Storage 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.

Registering multiple NetApp ONTAP providers through the Discover feature

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 [“Adding a site”](#) on page 24.
- When registering providers it is recommended to assign all related sources, such as hosting vCenters and related storage systems, to the same Site.
- For application server it is recommended to keep all resources, such as hosting vCenters and related storage systems, configured in the same Site.

Procedure

To register a NetApp ONTAP provider, complete the following steps:

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

Site

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

Name

A user-defined name for the 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 TLS** check box to enable an encrypted Secure Socket Layer connection. The typical default port is 80 for non TLS connections or 443 for TLS connections.

Comment

Optional provider description.

Run Inventory job after registration

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

Credentials

Select or create your NetApp ONTAP credentials. See [“Identities overview”](#) on page 67.

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

5. Click **OK**. IBM Storage 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.

Registering a Pure Storage FlashArray provider

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 [“Adding a site”](#) on page 24.
- When registering providers it is recommended to assign all related sources, such as hosting vCenters and related storage systems, to the same Site.
- For application server it is recommended to keep all resources, such as hosting vCenters and related storage systems, configured in the same Site.

Procedure

To register a Pure Storage FlashArray provider, complete the following steps:

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

Site

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

Name

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

Host Address

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

Port

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

Credentials

Select or create your Pure credentials. See [“Identities overview” on page 67](#).

5. Click **OK**. IBM Storage 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.

Registering a Security Scan Server in IBM Storage Copy Data Management

This procedure describes how to register a Security Scan Server in IBM Storage Copy Data Management.

Before you begin

- Register an IBM Storage Virtualize provider for steps refer to [Registering an IBM Storage Virtualize provider](#).
- If it is a virtual machine, register a VMware provider for steps refer to [Registering a VMware provider](#).
- If it is a physical machine, register an application server that follows:
 - InterSystems Database for steps refer to [Registering an Application Server - InterSystems Database](#)
 - SAP HANA for steps refer to [“Registering an Application Server - SAP HANA” on page 29](#).

Restriction:

- When a Security Scan Server is added to IBM Storage Copy Data Management, an index is created on the Security Scan Server. At any given time, only one index can be active. As a result, a Security Scan Server should be registered to no more than one IBM Storage Copy Data Management server.
- When a Security Scan Server option is selected for a backup job, make sure that the backup job consists of no more than one host as the backup target. There is no way to differentiate between the scan results of multiple hosts contained in the same backup job.
- To use the **Security Scan** feature, you must have installed the IBM Storage Sentinel Security Scan software with the license. For more information, refer to [“Installing IBM Storage Sentinel Security Scan” on page 18](#).

Procedure

To register a Security Scan Server, complete the following steps:

1. Log in to IBM Storage Copy Data Management. For more information, see [Access and Default Credentials](#).
2. Click the **Configure** tab. On the Views pane, select **Sites & Providers**.
3. In the Provider Browser pane, select **Security Scan Server**.
4. Right-click **Security Scan Server**. Then click **Register**. The Register dialog opens.
5. Populate the fields in the dialog:

Site

A user-defined provider location, created in the Sites & Providers view on the Configure tab. Select a site to which the Security Scan Server is to be assigned.

Name

Enter a name for the Security Scan Server. A user-defined name for the IBM provider. This name 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. Enter the IPv4 address of the Security Scan Server.

Type

Select the deployment type of the Security Scan Server. Select either Virtual or Physical.

vCenter

Select a vCenter if virtual is selected as the type.

Re-create index

The **Re-create index** should normally not be selected when you edit the **Security Scan Server**. The exceptions are described below.

When an index on the **Security Scan Server** is lost, you can re-create the index by selecting **Re-create index** and saving the **Security Scan Server Provider**.

Remember: If a member is moved from one federation to another, you must edit the member's existing registration, select **Re-create index**, and then save the changes in the **Edit Server Registration**.

Security Scan API Credential

Security Scan API Credential controls the API interfacing, which is same credential as used to log in web UI of Security Scan Server. Click **Select** or **New**. Security Scan API Credential allows the scan API credential to be able to register with the Security Scan Server that you can use for your UI.

System Credential

Select or create your credentials. System Credential is the SSH credentials for the Security Scan Server.

6. Click **Ok**.

Note: Registering Security Scan Server may take a few minutes to complete.

Results

Now you are ready to include the Security Scan Servers in an SLA policy for Safeguard for the backups.

What to do next

Create a Safeguarded Copy SLA policy for steps refer to [Creating a Safeguarded Copy SLA policy](#).

Registering an SMTP provider

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 [“Adding a site” on page 24](#).
- When registering providers it is recommended to assign all related sources, such as hosting vCenters and related storage systems, to the same Site.
- For application server it is recommended to keep all resources, such as hosting vCenters and related storage systems, configured in the same Site.
- Set up an SMTP server for email communications from IBM Storage Copy Data Management. Note that the SMTP provider is not cataloged.

Procedure

To register an SMTP provider, complete the following steps:

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

Name

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

Host Address

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

Port

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

Credentials

Select or create SMTP credentials. See [“Identities overview”](#) on page 67.

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

Comment

Optional provider description.

5. Click **OK**. IBM Storage 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.

Registering a VMware provider

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 [“Adding a site”](#) on page 24.
- When registering providers it is recommended to assign all related sources, such as hosting vCenters and related storage systems, to the same Site.
- For application server it is recommended to keep all resources, such as hosting vCenters and related storage systems, configured in the same Site.

Procedure

To register a VMware provider, complete the following steps:

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

Site

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

Name

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

Host Address

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

Port

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

Credentials

Select or create your VMware credentials. See [“Identities overview” on page 67](#).

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

Comment

Optional provider description.

Run Inventory job after registration

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

5. Click **OK**. IBM Storage 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 [“Adding credentials to a virtual machine” on page 45](#).

Registering multiple VMware providers through the Discover feature

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 [“Adding a site” on page 24](#).
- When registering providers it is recommended to assign all related sources, such as hosting vCenters and related storage systems, to the same Site.
- For application server it is recommended to keep all resources, such as hosting vCenters and related storage systems, configured in the same Site.
- 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.

Procedure

To register multiple VMware providers through the Discover feature, complete the following steps:

1. Click the **Configure** tab. On the Views pane, select **Sites & Providers**, then select the **Providers** tab.
2. In the Provider Browser pane, select **VMware**.
3. Right-click **VMware**. Then click **Discover**. The Discover VMware Providers dialog opens.
4. Enter an IP address or range of IP addresses associated with your VMware providers in the **IP Address** field.
5. Click **Discover**. Discovered VMware providers display.

6. Select providers to register along with universal custom parameters such as Site and credentials. To select individual parameters for each provider, click the parameters in the Custom Parameters field. Select specific Sites, credentials, ports and SSL parameters for each provider. If **Run Inventory job after registration** is selected, IBM Storage Copy Data Management creates a high-level Inventory job and automatically catalogs the objects on the provider.

7. Click **Register**. IBM Storage 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 [“Adding credentials to a virtual machine” on page 45](#).

What to do next

- Once providers are available in IBM Storage Copy Data Management and associated with a site, assign them to a resource pool. See [“Configure resource pools” on page 52](#).
- Add credentials to virtual machines in a VMware environment. See [“Adding credentials to a virtual machine” on page 45](#).

Related concepts

[“Sites and providers overview” on page 23](#)

A site is a user-defined grouping of providers that is generally based on location to help quickly identify and interact with data that is created through Copy Data Management jobs. Sites are assigned when registering providers. When creating Backup and Restore jobs, sites clearly identify where your data is replicated by location. Providers are physical servers that host objects and attributes. Once a provider is registered in IBM Storage Copy Data Management, cataloging, searching, and reporting can be performed.

[“Viewing a provider” on page 40](#)

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.

Related tasks

[“Editing a provider” on page 43](#)

Revise the properties of a provider as needed.

[“Unregistering a provider” on page 44](#)

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

[“Uploading an SSL Certificate” on page 330](#)

To establish secure connections in IBM Storage Copy Data Management, you must upload an SSL certificate through the web-based management console of the virtual machine where IBM Storage Copy Data Management is deployed.

Related reference

[“LDAP username syntax” on page 341](#)

LDAP form is used when setting up LDAP authentication in IBM Storage Copy Data Management.

Viewing 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.

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.

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.

Viewing a list of providers

This procedure describes how to view a list of providers.

Before you begin

Register providers before viewing them. See [“Registering a provider” on page 25](#).

Procedure

To view a list of providers, complete the following steps:

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

What to do next

Run an Inventory job against storage providers that have not been recently cataloged. See [“Jobs overview” on page 91](#).

Related concepts

[“Sites and providers overview” on page 23](#)

A site is a user-defined grouping of providers that is generally based on location to help quickly identify and interact with data that is created through Copy Data Management jobs. Sites are assigned when registering providers. When creating Backup and Restore jobs, sites clearly identify where your data is replicated by location. Providers are physical servers that host objects and attributes. Once a provider is registered in IBM Storage Copy Data Management, cataloging, searching, and reporting can be performed.

[“Registering a provider” on page 25](#)

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

Related tasks

[“Editing a provider” on page 43](#)

Revise the properties of a provider as needed.

[“Unregistering a provider” on page 44](#)

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

Viewing job definitions associated with providers

This procedure describes how to view job definitions associated with providers

Before you begin

Register providers before viewing them. See [“Registering a provider” on page 25](#).

Procedure

To view job definitions associated with providers, complete the following steps:

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

What to do next

Run an Inventory job against storage providers that have not been recently cataloged. See [“Jobs overview” on page 91](#).

Related concepts

[“Sites and providers overview” on page 23](#)

A site is a user-defined grouping of providers that is generally based on location to help quickly identify and interact with data that is created through Copy Data Management jobs. Sites are assigned when registering providers. When creating Backup and Restore jobs, sites clearly identify where your data is replicated by location. Providers are physical servers that host objects and attributes. Once a provider is registered in IBM Storage Copy Data Management, cataloging, searching, and reporting can be performed.

[“Registering a provider” on page 25](#)

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

Related tasks

[“Editing a provider” on page 43](#)

Revise the properties of a provider as needed.

[“Unregistering a provider” on page 44](#)

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

Editing a provider

Revise the properties of a provider as needed.

Before you begin

- Review the properties of your current providers. See **Sites and Providers Overview** for a list of supported providers.

Procedure

To edit the properties of a provider, complete the following steps:

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

Name

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

Host Address

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

Port

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

Credentials

Select or create your credentials. See [“Identities overview” on page 67](#).

Comment

Optional provider description.

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

What to do next

If the storage provider you edited has not recently been cataloged, catalog it. See [“Jobs overview” on page 91](#).

Related concepts

[“Sites and providers overview” on page 23](#)

A site is a user-defined grouping of providers that is generally based on location to help quickly identify and interact with data that is created through Copy Data Management jobs. Sites are assigned when registering providers. When creating Backup and Restore jobs, sites clearly identify where your data is replicated by location. Providers are physical servers that host objects and attributes. Once a provider is registered in IBM Storage Copy Data Management, cataloging, searching, and reporting can be performed.

[“Registering a provider” on page 25](#)

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

[“Viewing a provider” on page 40](#)

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.

Related tasks

[“Unregistering a provider” on page 44](#)

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

Unregistering a provider

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

Before you begin

Note: Ensure the provider you unregister is not associated with any defined job. To view job definitions associated with a provider, click the **Configure** tab.

- On the Views pane, select **Sites & Providers**, then select the **Providers** tab. Right-click a provider from the Provider Browser view, then select **View Relationship**. All associated job definitions that interact with the provider display.
- Review the properties of the provider to determine if you want to unregister it. See [“Editing a provider” on page 43](#).

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

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

Procedure

To unregister a provider, complete the following steps:

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

[“Sites and providers overview” on page 23](#)

A site is a user-defined grouping of providers that is generally based on location to help quickly identify and interact with data that is created through Copy Data Management jobs. Sites are assigned when registering providers. When creating Backup and Restore jobs, sites clearly identify where your data is replicated by location. Providers are physical servers that host objects and attributes. Once a provider is registered in IBM Storage Copy Data Management, cataloging, searching, and reporting can be performed.

[“Registering a provider” on page 25](#)

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

[“Viewing a provider” on page 40](#)

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.

Related tasks

[“Editing a provider” on page 43](#)

Revise the properties of a provider as needed.

Adding credentials to a virtual machine

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

Adding credentials for a single virtual machine

Before you begin

At least one VMware provider must be registered and cataloged in IBM Storage Copy Data Management. See [“Registering a provider” on page 25](#), and [“Creating a Database Inventory job definition” on page 97](#).

Procedure

To add credentials for a single virtual machine, complete the following steps:

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

Related concepts

[“Registering a provider” on page 25](#)

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

Related tasks

[“Creating a Database Inventory job definition” on page 97](#)

A Database Inventory job provides the framework to catalog and collect information about your application database servers.

Adding credentials for multiple virtual machines

Before you begin

At least one VMware provider must be registered and cataloged in IBM Storage Copy Data Management. See [“Registering a provider” on page 25](#), and [“Creating a Database Inventory job definition” on page 97](#).

Procedure

To add credentials for multiple virtual machines, complete the following steps:

1. Click the **Configure** tab. On the Views pane, select **Sites & Providers**, then select the **Providers** tab.
2. In the Provider Browser pane, expand the **VMware** object and right-click a **VMware provider**. Then click **Manage VMs**. The Manage Virtual Machines dialog opens.

3. Enter a wildcard to search for virtual machines available on the VMware provider. For example, vm* or vm[1-50].
4. Select virtual machines with universal credentials.
5. Enter the universal credential information for the virtual machines, along with the credential type and instance name if applicable.
6. To apply System credentials to application instances (for example, SQL instances), enable the **Use System Credentials for apps** option. Note that System credentials are always required. If your application instances use credentials that differ from your System credentials, you must repeat the above procedure for each application instance.
7. Click **Close** to exit the Manage Virtual Machines dialog.

What to do next

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 concepts

[“Registering a provider” on page 25](#)

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

Related tasks

[“Creating a Database Inventory job definition” on page 97](#)

A Database Inventory job provides the framework to catalog and collect information about your application database servers.

Chapter 7. Requiring multifactor authentication

Starting with IBM Storage Copy Data Management 2.2.21, you can set up multifactor authentication (MFA) on IBM Storage Copy Data Management new and existing user accounts. MFA provides an extra layer of protection by requiring users to use a password and a time-based one-time password (TOTP) to sign in. The implementation covers enabling or disabling MFA, user enrollment, login with MFA, token expiration, and administrative control over MFA settings.

A user with ADMIN role can enable, disable and expire multifactor authentication (MFA) on both new and existing accounts. Multifactor authentication requires users to verify their identity by using more than one method. When you require a multifactor authentication, you must provide a one-time passcode in addition to the traditional password. The passcode is valid only for the current session and is generated on a trusted device. A trusted device is a device that only the user can access. To use this feature, you must install a security application on your trusted devices. The trusted device is typically a mobile, but can be a different device such as a tablet or laptop. The security application generates a time-based one-time password (TOTP) that is used during the sign-in process.

A regular user can only enroll and login by using MFA.

The following considerations apply to configure the time-based one-time (TOTP) multifactor authentication:

- By default, time-based one-time (TOTP) is disabled. If you want to enable to the time-based one-time (TOTP) for a user account, you must manually enable it in the IBM Storage Copy Data Management user interface.
- A user with ADMIN role has permissions or access rights to enable or disable the time-based one-time (TOTP) multifactor authentication (MFA).
- A user with ADMIN role has permissions or access rights to expire the existing MFA TOTP secret keys.
- The time that is displayed on the user's mobile device or workstation, where the security application is installed, must be in synchronization with the IBM Storage Copy Data Management server time. You must use NTP server to ensure the time sync between IBM Storage Copy Data Management server and your mobile device or workstation.
- When time-based one-time (TOTP) is enabled for a user, all existing sessions that are associated with that user expire.

After your account is successfully set up with MFA, the account can be accessed only by specifying the password and a TOTP passcode. The additional layer of protection authorizes only the rightful owner to access the account.

Multifactor authentication for the ADMIN user

You can enable time-based one-time (TOTP) multifactor authentication (MFA) on the IBM Storage Copy Data Management ADMIN account.

Procedure

To enable time-based one-time (TOTP) multifactor authentication on the IBM Storage Copy Data Management ADMIN account, complete the following steps:

1. Login as the ADMIN.
2. Select the **ADMIN** account.
3. Select **enable TOTP** from **MFA Actions**.
4. Click **Yes** in the pop up confirmation.

Note: After enabling TOTP for the admin user, admin must enroll by using the IBM Storage Copy Data Management portal. Admin user must use the same TOTP token to log-in to the **Admin** console .

Creating multifactor authentication user account

You can create a multifactor authentication (MFA) user account on IBM Storage Copy Data Management.

Procedure

To create multifactor authentication user account on IBM Storage Copy Data Management, complete the following steps:

1. Login as the ADMIN.
2. In the navigation panel, click **Access Control > New > Create Native User**.
3. Update the **Username** and **Password**.
4. Select the **Resource Pools**, and select the **roles/ permissions for the resource pool**.
5. Click **Finish**.

Enabling time-based one-time multifactor authentication for a user

You can enable time-based one-time (TOTP) multifactor authentication (MFA) on the IBM Storage Copy Data Management user accounts. By default, time-based one-time (TOTP) is disabled for all user accounts and must be enabled on a per-user basis by the ADMIN user.

Procedure

To enable time-based one-time (TOTP) multifactor authentication on a IBM Storage Copy Data Management user account, complete the following steps:

1. Login as the ADMIN.
2. Select the user account.
3. Select **enable TOTP** from **MFA Actions**.
4. Click **Yes** in the pop up confirmation.

Setting up a multifactor authentication user account in the browser

You can register a multifactor authentication (MFA) IBM Storage Copy Data Management user account in the browser. You must install a security application in your mobile device or workstation.

Procedure

To register a multifactor authentication IBM Storage Copy Data Management user account on the browser, complete the following steps:

1. Sign on to the IBM Storage Copy Data Management using the user ID and password of MFA-enabled user. The IBM Storage Copy Data Management displays a QR code that encodes a shared secret.
2. Scan the QR code by using the security application that generates a time-based one-time password (TOTP) setup on the user's mobile device or workstation.

Note: Some of the supported security applications are as follows:

- IBM Verify
 - Duo Mobile
 - Google Authenticate
 - Microsoft Authenticator
3. Enter the TOTP passcode that is generated by the security application in the Passcode field.
 4. Click **Continue** to complete the MFA set up.

Expiring time-based one-time secret key of a user account

A user with ADMIN role has permission to expire time-based one-time (TOTP) secret keys for the IBM Storage Copy Data Management user accounts that are enabled with multifactor authentication (MFA).

Procedure

To expire time-based one-time (TOTP) secret key of a user account that is enabled with multifactor authentication (MFA), complete the following steps:

1. Log in as the ADMIN.
2. Select the user and click **MFA Accounts**.
3. Select **Expire TOTP Secret**.

Disabling multifactor authentication

A user with ADMIN role has permission to disable time-based one-time (TOTP) of the IBM Storage Copy Data Management user accounts that are enabled with multifactor authentication (MFA).

Procedure

To disable time-based one-time (TOTP) multifactor authentication (MFA) for a IBM Storage Copy Data Management user account, complete the following steps:

1. Log in as the ADMIN.
2. Select the user and then click **MFA Actions**.
3. Select **Disable TOTP**.
4. Select **Yes** in the pop up confirmation.

Chapter 8. Role-based access control

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

Role-based access control overview

Role-based access control allows you to set the resources and permissions available to IBM Storage Copy Data Management accounts.

Through role-based access control you can tailor IBM Storage 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 Storage 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 Storage Copy Data Management or imported as part of an LDAP group.

- [Configure Resource Pools](#)
- [Configure Roles](#)
- [Configure Accounts](#)
- [Configure Tenants](#)

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

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

Resource Pools

A resource pool defines the resources that will be made available to an account. Every provider added to IBM Storage Copy Data Management, such as storage devices and LDAP servers, can be included in a resource pool, along with individual IBM Storage 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 Storage 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](#).

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 Storage 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 Storage Copy Data Management Backup and Restore jobs, the role will determine how a user can interact with the jobs. Permissions can be set to allow a user to create, view, and run the Backup and Restore jobs defined in a resource pool, but not delete them. Similarly, permissions can be set to create administrator accounts, allowing a user to create and edit other accounts, set up sites and resources, and interact with all of the available IBM Storage Copy Data Management features. See [Configure Roles](#).

Accounts

An account associates a resource pool with a role. To enable a user to log on to IBM Storage Copy Data Management and use its functions, you must first add the user to IBM Storage 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](#).

Related concepts

[“Configure resource pools” on page 52](#)

[“Configure roles” on page 54](#)

[“Configure accounts” on page 55](#)

[“VMware admin role-based access control configuration” on page 57](#)

[“NetApp ONTAP admin role-based access control configuration” on page 59](#)

[“IBM admin role-based access control configuration” on page 60](#)

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 Storage Copy Data Management, such as storage devices and LDAP servers, can be included in a resource pool, along with individual IBM Storage 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 Storage Copy Data Management search and reporting functionality. When the resource pool is associated with a role and an account, the account user will only see the screens associated with search and reporting, and will only have access to the storage devices defined in the resource pool.

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

Review the following consideration for configuring resource pools.

- 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](#).
- Add providers to IBM Storage Copy Data Management and associate them with a site. See [Register a Provider](#).

Adding a resource pool

A resource pool is a component of the role-based access system, and defines the resources that will be made available to an account.

Procedure

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

5. Click the **Step 3: Reports** tab. Select one or more report types or individual reports to include in the resource pool.
6. Click the **Step 4: Applications** tab. Select one or more application servers to include in the resource pool.
7. Click the **Step 5: Identities** tab. Select one or more keys and credentials to include in the resource pool.
8. Click the **Step 6: Access Control** tab. Select security options that will be configurable by accounts associated with this resource pool. Available options include All Roles, All Accounts, All Resource Pools, and All SLA Policies. For example, if **All Resource Pools** is selected in this step, users associated with this resource pool can create, view, edit, and delete Resource Pools, if paired with the necessary "resourcepool" permission, set on the **Roles** panel.
9. Click the **Step 7: Screens** tab. Select the IBM Storage Copy Data Management screens to include in the resource pool.
10. Click the **Step 8: Finish** tab. Enter a name for your resource pool and a meaningful description. When you are satisfied that the entered information is correct, click **Finish**. The resource pool appears on the All Resources pane and can be applied to new and existing accounts.

Editing a resource pool

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

Procedure

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

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

Deleting a resource pool

Delete a resource pool when it becomes obsolete.

About this task

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

Remember: The All Resources resource pool cannot be deleted.

Procedure

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

What to do next

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](#).

Related concepts

[“Configure roles” on page 54](#)

[“Configure accounts” on page 55](#)

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 Storage 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 Storage Copy Data Management Backup and Restore jobs, the role will determine how a user can interact with the jobs. Permissions can be set to allow a user to create, view, and run the Backup and Restore jobs defined in a resource pool, but not delete them. Similarly, permissions can be set to create administrator accounts, allowing a user to create and edit other accounts, set up sites and providers, and interact with all of the available IBM Storage Copy Data Management features.

Review the following consideration for configuring resource pools.

- 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](#).
- Add providers to IBM Storage Copy Data Management and associate them with a site. See [Register a Provider](#).
- Once providers are available in IBM Storage Copy Data Management and associated with a site, assign them to a resource pool. See [Configure Resource Pools](#).

Adding a role

The role appears on the All Roles pane and can be applied to new and existing accounts.

Procedure

1. Click the **Configure** tab. On the Views pane, select **Access Control**, then select the **Roles** tab.
2. In the All Roles pane, click **New**. The New Role dialog opens.
3. Enter a role name and a meaningful description.
4. Select IBM Storage Copy Data Management features to add to the role, such as reports, jobs, and sites as well as provider types, such as VMware, LDAP, and SMTP.
5. When a feature is added to the role, it displays in the Permissions pane. Select permissions for the feature.

For example, if the Site feature is added to the role, the following Site-based permissions are available: Create, View, Edit, Delete and All Permissions. If the Delete permission is excluded from the role, accounts associated with this role can create, view, and edit Sites, but cannot delete them.

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

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

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

Permissions are then added to the selected roles.

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

Editing a role

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

About this task

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

Remember: The SYSADMIN and USER roles cannot be edited.

Procedure

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

Deleting a role

Delete a resource pool when it becomes obsolete.

About this task

Delete a role when it becomes obsolete.

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

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

Procedure

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

What to do next

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

Related concepts

[“Configure resource pools” on page 52](#)

[“Configure accounts” on page 55](#)

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 Storage Copy Data Management and use its functions, you must first add the user to IBM Storage 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.

Review the following consideration for configuring resource pools.

- 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](#).
- Add providers to IBM Storage Copy Data Management and associate them with a site. See [Register a Provider](#).
- Once providers are available in IBM Storage Copy Data Management and associated with a site, assign them to a resource pool. See [Configure Resource Pools](#).
- 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](#).

Adding a native account to IBM Storage Copy Data Management

An account is a component of the role-based access system, and is used to associate resource pools and roles with a user.

Procedure

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

Importing LDAP groups into IBM Storage Copy Data Management

You can import LDAP groups into IBM Storage Copy Data Management.

Procedure

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

Editing an account

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

Procedure

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

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

Deleting an account

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

Procedure

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

What to do next

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

Related concepts

[“Configure resource pools” on page 52](#)

[“Configure roles” on page 54](#)

VMware admin role-based access control configuration

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

Resource pool configuration

<i>Table 5. Resource pool configuration</i>	
Resource pool configuration	Steps
Providers Tab	<ol style="list-style-type: none"> 1. Set up the Providers screen to include the root level for all VMware resources and the IBM Storage Virtualize resources. 2. Select a specific SMTP server.
Jobs Tab	<ol style="list-style-type: none"> 1. Select the root level of all VMware related jobs. 2. Select the root Report job. 3. Select the root All SLA Policies. 4. Select the root All Schedules.
Reports Tab	<ol style="list-style-type: none"> 1. Under Protection Compliance, select the VMware related reports. 2. Under Storage Utilization, select the VMware related reports.
Access Control Tab	<ol style="list-style-type: none"> 1. No Security Resources will be assigned to this role.

<i>Table 5. Resource pool configuration (continued)</i>	
Resource pool configuration	Steps
Screens Tab	<ol style="list-style-type: none"> 1. Select all available Screens except Logs. The Logs function contains audit logs that this user should not have access to. 2. Name and submit the Resource Pool.
Role Configuration	<ol style="list-style-type: none"> 1. Select the permissions listed above. 2. Name and submit the Role.
VMware Admin Account Configuration	<ol style="list-style-type: none"> 1. Create a new account and link it to the newly created VMware Admin Resource Pool and Role.

Related concepts

[“Configure resource pools” on page 52](#)

[“Configure roles” on page 54](#)

[“Configure accounts” on page 55](#)

Database admin role-based access control configuration

The Database admin (DBA) role is applicable for any database that IBM Storage Copy Data Management supports. The role manages the database resource and runs Backup and Restore jobs on these database resources. The DBA also has permissions to configure work flow templates to allow other IBM Storage Copy Data Management users to run Backup and Restore jobs by using database resources without having direct access to them.

Resource pool configuration

<i>Table 6. Resource pool configuration</i>	
Resource pool configuration	Steps
Providers Tab	<ol style="list-style-type: none"> 1. Set up the Providers to include the root level for all database resources and the storage resources. If database server is virtual, include the VMware resource. 2. Select a specific SMTP server, if applicable.
Jobs Tab	<ol style="list-style-type: none"> 1. Select the root level of all database-related jobs. 2. Select the root All SLA Policies. 3. Select the root All Schedules.
Applications Tab	<ul style="list-style-type: none"> • Select the root level of all the database resources.
Reports Tab	<ul style="list-style-type: none"> • No Reports resources are assigned to this role.
Identities Tab	<ul style="list-style-type: none"> • Select all the keys or credentials for the database resources.
Access Control Tab	<ul style="list-style-type: none"> • No Security Resources are assigned to this role.
Screens Tab	<ul style="list-style-type: none"> • Select all available Screens and exclude Configure, Marketplace, and Logs. The DBA should not be able to configure the providers. The logs contain audit logs that the user should not have access.

Table 6. Resource pool configuration (continued)

Finish Tab	<ol style="list-style-type: none"> 1. Provide a name for the resource pool. 2. Submit the Resource Pool.
------------	--

Role Configuration

1. Select the permissions required for each resource.
2. Provide a name to the role and click **Finish** to save the role.

Database Admin Account Configuration

Create a new account and link it to the newly created DBA Resource Pool and Role.

NetApp ONTAP admin role-based access control configuration

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

Resource pool configuration

Table 7. Resource pool configuration	
Resource pool configuration	Steps
Providers Tab	<ol style="list-style-type: none"> 1. Set up the Providers screen to include the root level for all NetApp resources. 2. Select a specific SMTP server.
Jobs Tab	<ol style="list-style-type: none"> 1. Select the root level of all NetApp related jobs. 2. Select the root Report job. 3. Select the root All SLA Policies. 4. Select the root All Schedules.
Reports Tab	<ol style="list-style-type: none"> 1. Select the root File Analytics tree. 2. Under Protection Compliance, select the NetApp RPO and NetApp Protection Usage. 3. Select the root Storage Protection tree. 4. Under Storage Utilization, select the NetApp related reports.
Access Control Tab	<ol style="list-style-type: none"> 1. No Security Resources will be assigned to this role.
Screens Tab	<ol style="list-style-type: none"> 1. Select all available Screens except Logs. The Logs function contains audit logs that this user should not have access to. 2. Name and submit the Resource Pool.
Role Configuration	<ol style="list-style-type: none"> 1. Select the permissions listed above. 2. Name and submit the Role.

Table 7. Resource pool configuration (continued)

Resource pool configuration	Steps
NetApp Admin Account Configuration	1. Create a new account and link it to the newly created VMware Admin Resource Pool and Role.

Related concepts

[“Configure resource pools” on page 52](#)

[“Configure roles” on page 54](#)

[“Configure accounts” on page 55](#)

IBM admin role-based access control configuration

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

Resource pool configuration

Table 8. Resource pool configuration

Resource pool configuration	Steps
Providers Tab	<ol style="list-style-type: none"> 1. Set up the Providers screen to include the root level for all IBM storage resources. 2. Select a specific SMTP server.
Jobs Tab	<ol style="list-style-type: none"> 1. Select the root level of IBM Inventory, IBM Backup, and IBM Restore jobs. 2. Select the root Report job. 3. Select the root All SLA Policies. 4. Select the root All Schedules.
Reports Tab	<ol style="list-style-type: none"> 1. Under Protection Compliance, select the IBM RPO Compliance report. 2. Under Storage Utilization, select the IBM related reports.
Access Control Tab	<ol style="list-style-type: none"> 1. No Security Resources will be assigned to this role.
Screens Tab	<ol style="list-style-type: none"> 1. Select all available Screens except Logs. The Logs function contains audit logs that this user should not have access to. 2. Name and submit the Resource Pool.
Role Configuration	<ol style="list-style-type: none"> 1. Select the permissions listed above. 2. Name and submit the Role.
IBM Admin Account Configuration	<ol style="list-style-type: none"> 1. Create a new account and link it to the newly created IBM Admin Resource Pool and Role.

Related concepts

[“Configure resource pools” on page 52](#)

[“Configure roles” on page 54](#)

[“Configure accounts” on page 55](#)

Pure Storage FlashArray admin role-based access control configuration

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

Resource pool configuration

Resource pool configuration	Steps
Providers Tab	<ol style="list-style-type: none">1. Set up the Providers screen to include the root level for all Pure Storage resources.2. Select a specific SMTP server.
Jobs Tab	<ol style="list-style-type: none">1. Select the root level of all Pure Storage related jobs.2. Select the root Report job.3. Select the root All SLA Policies.4. Select the root All Schedules.
Reports Tab	<ol style="list-style-type: none">1. Under Protection Compliance, select the Pure Storage FlashArray RPO Compliance.2. Under Storage Utilization, select the Pure Storage related reports.
Access Control Tab	<ol style="list-style-type: none">1. No Security Resources will be assigned to this role.
Screens Tab	<ol style="list-style-type: none">1. Select all available Screens except Logs. The Logs function contains audit logs that this user should not have access to.2. Name and submit the Resource Pool.
Role Configuration	<ol style="list-style-type: none">1. Select the permissions listed above.2. Name and submit the Role.
Pure Storage Admin Account Configuration	<ol style="list-style-type: none">1. Create a new account and link it to the newly created Pure Storage Admin Resource Pool and Role.

Related concepts

[“Configure resource pools” on page 52](#)

[“Configure roles” on page 54](#)

[“Configure accounts” on page 55](#)

Chapter 9. Configure tenants

A tenant is a grouping of resources and users that are administered by a tenant administrator. An IBM Storage 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 Storage 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 Storage Copy Data Management resources and actions available within a tenant. A built-in Tenant role may be selected, which gives tenant users the ability to register resources, create job definitions, and other predefined IBM Storage 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 Storage Copy Data Management username field: tenant1/tenant_user.

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

Review the following consideration for configuring tenants.

- Review Best Practices for [Configuring Tenants](#).
- 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](#).
- 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. See [Configure Roles](#).

Adding a tenant

A tenant is a grouping of resources and users that are administered by a tenant administrator.

Procedure

1. Click the **Configure** tab. On the Views pane, select **Access Control**, then select the **Tenants** tab.
2. Click **New**. The New Tenant editor opens.
3. In the **Enter Tenant Info** section, enter a name for the tenant in the Tenant Name field as well as a tenant administrator name in the **Tenant Admin Name** field. Enter and confirm a password for the tenant administrator.
4. In the **Select resource pools** section, select one or more resource pools to add to the tenant.
5. In the **Select the roles/permissions for the resource pool** section, click **Click to select roles** to assign roles to the selected resource pools. Note that a built-in Tenant role may be selected, which gives tenant users the ability to register resources, create job definitions, and other predefined IBM Storage 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.

Editing a tenant

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

Procedure

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

Deleting a tenant

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

Procedure

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

Related concepts

[“Best practices for configuring tenants” on page 65](#)

[“Configure resource pools” on page 52](#)

Chapter 10. Best practices for configuring tenants

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

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

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

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

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

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

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

General recommendations

- For a tenant admin, it is recommended to create two resource pools. In the first resource pool, add the providers to be made available to the tenant, and assign a Read Only permission to the resource pool. In the second resource pool, assign jobs, security, and screens, and assign the Create Only permission. Once complete, assign both resource pools to the tenant.
- When configuring a resource pool for a tenant user, it is recommended to exclude the security resources found in the **Step 5: Security** step or the Home and Logs resources found in the **6. Screens** step. These resources contain general IBM Storage Copy Data Management information that may not apply to the tenant user.
- The built-in All Resources resource pool should not be assigned to a tenant as it includes all of the resources in the IBM Storage 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 concepts

[“Configure tenants” on page 63](#)

[“Configure resource pools” on page 52](#)

Chapter 11. Identities

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

Identities overview

Credentials and keys are required to access some providers in IBM Storage Copy Data Management. Credentials and keys are configured through the **Identities** view.

Some features in IBM Storage Copy Data Management require credentials and keys to access your providers. For example, IBM Storage 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 Storage 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 Storage Copy Data Management needs to perform its tasks.

Related concepts

[“Adding a key” on page 67](#)

Credentials and keys may be required to access providers added to IBM Storage Copy Data Management.

Related tasks

[“Adding a credential” on page 70](#)

The procedures describe how to add a credential.

Adding a key

Credentials and keys may be required to access providers added to IBM Storage Copy Data Management.

For example, IBM Storage 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 Storage 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 Storage Copy Data Management needs to perform its tasks.

IBM Storage 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](#).

Adding an SSH key through the Generate a keypair for me method and registering an associated provider

The procedures describe how to add an SSH key through the **Generate a key pair for me** method and registering an associated provider

Procedure

1. In IBM Storage 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 Storage Copy Data Management. Paste and save the generated public key to the `authorized_keys` file.
6. In IBM Storage Copy Data Management, click the **Configure** tab. On the Views pane, select **Sites & Providers** . The Provider Browser opens.
7. Right-click **Oracle** in the Provider Browser, then click **Register**. The Register Oracle Server dialog opens.
8. Select a Site, enter a Name and Host Address.
9. Select **Key** as the Authentication type. Enter the Oracle username, then select the key created in Step 2 in the Key field. Click **OK**.

Adding an SSH key through the I will provide a keypair method and registering an associated provider

The procedures describe how to add an SSH key through the **I will provide a keypair** method and registering an associated provider

About this task

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

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

Procedure

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

```
$ ssh-keygen
```

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

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

4. Supply a passphrase and press Enter. Otherwise, press Enter for no passphrase.
5. If a passphrase was supplied, enter it again. Press Enter.

6. The key generation will produce two files, one with the path name supplied in the previous steps for the private key, and another ending in `.pub` is the public key. Using the default naming, this will be `id_rsa` and `id_rsa.pub`. The generated public key (`ida_rsa.pub`) will need to be transferred to the server to which the IBM Storage Copy Data Management appliance will connect. In this example, it will be an Oracle server. Transfer the public key to the Oracle server. For the remainder of this procedure, it is assumed that the keypair is saved in the default location using the default file names for the keypair: `/home/<user_account>/.ssh/`. If the keypair is created using a different file name, use that file name in the steps that follow.
7. On the server to which the IBM Storage Copy Data Management appliance will connect and to which the public key has been copied, the key (`ida_rsa.pub`) will need to be appended to the user's `authorized_keys` file. The `authorized_keys` file is generally found in the user's SSH directory. For example, it may be found at the following location: `/home/<user_account>/.ssh/authorized_keys`. If the `authorized_keys` file does not exist, consult the operating system's documentation for the procedure to properly creating this file. If the file exists, append the contents of the public key to the `authorized_keys` file. If this is not being done from the account that contains the `authorized_keys` file, it may be necessary to enter the `su` command to switch to that user. The step below assumes that you are logged into the server with the account that contains the `authorized_keys` file:

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

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

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

Adding a credential

The procedures describe how to add a credential.

Procedure

1. Click the **Configure** tab. On the Views panel, 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** panel and can be applied to new and existing storage providers.

Related concepts

[“Identities overview” on page 67](#)

Credentials and keys are required to access some providers in IBM Storage Copy Data Management. Credentials and keys are configured through the **Identities** view.

[“Registering a provider” on page 25](#)

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

Editing a credential

Revise a credential to change the associated username and password.

Procedure

1. Click the **Configure** tab. On the Views panel, 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.

Related concepts

[“Identities overview” on page 67](#)

Credentials and keys are required to access some providers in IBM Storage Copy Data Management. Credentials and keys are configured through the **Identities** view.

[“Registering a provider” on page 25](#)

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

Deleting a credential

The procedures describe how to delete a credential.

Procedure

1. Click the **Configure** tab. On the Views panel, 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 concepts

[“Identities overview” on page 67](#)

Credentials and keys are required to access some providers in IBM Storage Copy Data Management. Credentials and keys are configured through the **Identities** view.

[“Registering a provider” on page 25](#)

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

Chapter 12. Configure SLA policies

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

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

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

Creating an SLA policy

The procedures describe how to create an SLA policy.

Before you begin

- Register storage providers to be used in the SLA Policies. Assign the storage providers to Sites. See [Register a Provider](#) and [Add a Site](#).
- Create an account with the necessary permissions to create and run Backup jobs. See [Role-Based Access Control Overview](#).
- All related NetApp ONTAP storage resources associated with a VMware provider must be added to IBM Storage Copy Data Management, which include NetApp ONTAP storage controllers and clusters. See [Register a Provider](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

- Note that a VADP-based VM Replication of a virtual machine with vRDM will convert the vRDM to a VMDK. When restoring the virtual machine, consider the size requirements of the virtual machine in addition to the vRDM when selecting a destination datastore. The original datastore may not have enough free space to store the converted vRDM file. Physical RDMs are not supported
- Note that VMware Backup and Restore jobs only support vCenters or ESX hosts running vSphere 6.0 through 7.0.
- Restoring to an older snapshot is not supported if a new snapshot is being used for mirror replication.

Considerations for VMware

- SLA Policies that include virtual machines stored on virtual volume (VVOL) datastores through VM Replication sub-policies are supported. Replication is supported on the VM Replication target.
- Storage Controller Volume 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.

About this task

Delete a job definition when it becomes obsolete. This keeps your operations current.

Procedure

1. Click the **Configure** tab. On the Views pane, select **SLA Policies**. The All SLA Policies pane opens.

2. In the All SLA Policies pane, click **New**. The New SLA Policies pane opens.
3. Select a type of policy to create based on your storage provider. Select **NetApp ONTAP** to create a NetApp ONTAP Backup job containing snapshots, VM copies, mirrors and vaults, or **IBM** to create an IBM Backup job containing FlashCopies, Global Mirrors with Change Volumes, and VM Copies. VMware Backup jobs support IBM and NetApp ONTAP SLA Policies depending on your storage provider.
4. Enter a name and a meaningful description of the SLA Policy.

What to do next

- Create a Backup job definition that utilizes an SLA Policy.

Related tasks

[“Creating an IBM Storage Virtualize Backup job definition” on page 131](#)

The procedures describe how to create an IBM Storage Virtualize Backup job definition.

[“Creating a NetApp ONTAP Backup job definition” on page 140](#)

The procedures describe how to create a NetApp ONTAP Backup job definition.

[“Creating a Pure Storage FlashArray Backup job definition” on page 142](#)

The procedures describe how to create a Pure Storage FlashArray Backup job definition.

[“Creating a VMware Backup job definition” on page 145](#)

The procedures describe how to create a VMware Backup job definition.

Configure IBM Storage Virtualize SLA policies

The procedures describe how to configure IBM Storage Virtualize SLA policies.

Adding a VM Replication sub-policy to an IBM Storage Virtualize SLA Policy

The procedures describe how to add a VM Replication sub-policy to an IBM Storage Virtualize SLA Policy.

Procedure

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Tip: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add VM Copy**.
3. In the Destination pane select an IBM host destination from the list of available resources as the VM Replication destination, along with an associated storage pool. If no storage pool is selected, the storage pool with the largest amount of space available is chosen by default. To select the original target destination, select **Use Original**.
4. In the Options pane set the VM Replication sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Target Volume Prefix Label

Enter an optional label to identify the target volume. This label is added as a prefix to the volume name created by the job.

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

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Tip: 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 Storage Copy Data Management. The default initial label is VM Copy0.

Protocol

If more than one storage protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

Full Copy Method

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

Destination storage limit in GB / Destination volumes limit

Specify quotas for storage usage and the number of volume created on the destination for all jobs utilizing the SLA Policy.

Adding a FlashCopy sub-policy to an IBM Storage Virtualize SLA Policy

The procedures describe how to add a FlashCopy® sub-policy to an IBM Storage Virtualize SLA Policy.

Procedure

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Tip: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. In the Target FlashCopy Storage Pool pane select an IBM host destination from the list of available resources as the FlashCopy destination, along with an associated storage pool. If no storage pool is selected, the storage pool with the largest amount of space available is chosen by default. To select the original target destination, select **Use Original**. Note that you can select multiple storage pools from multiple IBM resources, but only one storage pool is allowed for each node.
3. To enable an Incremental FlashCopy, select Enable Incremental FlashCopy in the Incremental FlashCopy Storage Pool pane. Select an IBM host destination from the list of available resources as the FlashCopy destination, along with an associated storage pool. If no storage pool is selected, the storage pool with the largest amount of space available is chosen by default. To select the original target destination, select Use Original.

If the Enable Incremental FlashCopy option is selected, note that the base FlashCopy will be sent to the destination selected in the Incremental FlashCopy Storage Pool pane. Subsequent incremental FlashCopies will be sent to the destination selected in the Target FlashCopy Storage Pool pane.

Note that the Target FlashCopy Storage Pool must reside on the same storage system as the Incremental FlashCopy Storage Pool.

Tip: When editing an existing SLA Policy, Incremental FlashCopy options cannot be altered.

4. Click **Add FlashCopy**. In the Options pane, set the FlashCopy sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Name

Enter an optional name to replace the default FlashCopy sub-policy name displayed in IBM Storage Copy Data Management. The default initial name is FlashCopy0.

FlashCopy Volume Prefix

Enter an optional label to identify the FlashCopy. This label is added as a prefix to the FlashCopy name created by the job.

Tip: FlashCopy labels must contain only alphanumeric characters and underscores.

Do not stretch FlashCopy (Enhanced Stretch Cluster Feature)

Enabling this option will not create stretch FlashCopies. This is only applies to volumes that are stretched in storages that have the Enhanced Stretch Cluster feature enabled.

Creating a Safeguarded Copy SLA policy

This procedure describes how to create a Safeguarded Copy (SGC) in IBM Storage Copy Data Management. The Safeguarded Copy function isolates backup copies from production data, so if a cyberattack occurs, you can quickly recover and restore data from Safeguarded copies.

Before you begin

- You must have a volume group that is created on the Flash System and must have assigned with at least one default safeguarded policy assigned. For steps refer to [Getting started with Safeguarded Copy function](#).
- If you are using IBM Storage Sentinel, you can scan snapshots for ransomware after you register a Security Scan Server by selecting Perform Security Scan. To register a Security Scan Server, see [Registering a Security Scan Server](#).

About this task

The Safeguarded Copy feature creates safeguarded backups that are not accessible by the host system and protects these backups from corruption that can occur in the production environment. You can define a Safeguarded Copy schedule to create multiple backups regularly, such as hourly or daily. You can also restore a backup to the source volume or to a different volume.

Procedure

1. Click the **Configure** tab. On the Views pane, select **Sites & Providers**, then expand the IBM Storage Virtualize and add the supported storage.
2. On the Views pane, select **SLA Policies**. The All SLA Policies pane opens.
3. In the All SLA Policies pane, click **New**. The New SLA Policies pane opens.
4. Add a sub-policy (SLA Policy) to an IBM Storage Virtualize SLA policy.
 - a) Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the Frequency field, select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.
 - b) Click **Add Safeguarded Copy**.
 - c) In the Associated Safeguarded Volume Group pane, expand the storage device and select the volume group that you want to be back up as Safeguarded Copy. Any volume group that you want to back up as Safeguarded Copy must be a volume that belongs to one of these groups. If it is not a member of any of these groups it will not back up as a Safeguarded Copy.

Note: The Associated Safeguarded Volume Group lists only those volume groups that have the Safeguarded Copy policy applied on the storage array side.
 - d) In the Options pane, set the Safeguarded 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.

Name

Enter an optional name to replace the default FlashCopy sub-policy name displayed in IBM Storage Copy Data Management. The default name is SaFeguarded Copy0.

FlashCopy Volume Prefix

Enter an optional label to identify the FlashCopy. This label is added as a prefix to the FlashCopy name created by the job.

Tip: FlashCopy labels must contain only alphanumeric characters and underscores.

Perform Security Scan

You must enable this and select your security scan servers. This allows to scan for every backup number you have specified.

- e) Enter a name for the new sub-policy (SLA Policy).
- f) Click **Finish**.

Adding a Global Mirror with Change Volumes sub-policy to an IBM Storage Virtualize SLA Policy

The procedures describe how to add a Global Mirror with Change Volumes sub-policy to an IBM Storage Virtualize SLA Policy.

Procedure

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Tip: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add Global Mirror with Change Volumes**.
3. In the Global Mirror with Change Volumes Destination pane select an IBM host destination from the list of available resources as the Global Mirror destination, along with an associated storage pool. If no storage pool is selected, the storage pool with the largest amount of space available is chosen by default. To select the original target destination, select **Use Original**.
4. In the Options pane set the Global Mirror with Change Volumes sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Name

Enter an optional name to replace the default Global Mirror sub-policy name displayed in IBM Storage 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 Storage Copy Data Management.

Volume Prefix Label

Enter an optional label to identify the volume. This label is added as a prefix to the volume name created by the job and cannot be edited after the job is submitted.

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

Cycle Period (seconds)

Specify the time in which the change volumes will be refreshed with a consistent copy of the data. If a copy does not complete in the cycle period, the next cycle period will not start until the copy is complete. The range of possible values is 60 through 86400. The default is 300.

Global Mirror Volume Prefix

Enter an optional label to identify the Global Mirror. This label is added as a prefix to the Global Mirror name created by the job.

Tip: Global Mirror labels must contain only alphanumeric characters and underscores.

Destination storage limit in GB / Destination volumes limit

Specify quotas for storage usage and the number of volume created on the destination for all jobs utilizing the SLA Policy.

Configure IBM Storage Virtualize for Snapshot SLA policies

The procedures describe how to configure IBM Storage Virtualize for Snapshot SLA policies.

Adding a Snapshot sub-policy to an IBM Storage Virtualize for Snapshot SLA Policy

The procedures describe how to add a Snapshot sub-policy to an IBM Storage Virtualize for Snapshot SLA Policy.

Procedure

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Tip: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. In the Target Snapshot Storage Pool pane select an IBM host destination from the list of available resources as the Snapshot 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. 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 Storage Copy Data Management. The default initial name is Snapshot0.

Snapshot Prefix Label

Enter an optional label to identify the Snapshot. This label is added as a prefix to the Snapshot name created by the job.

Tip: Snapshot Prefix Label labels must contain only alphanumeric characters and underscores.

Creating a Safeguarded Copy SLA policy

This procedure describes how to create a Safeguarded Copy (SGC) in IBM Storage Copy Data Management. The Safeguarded Copy function isolates backup copies from production data, so if a cyberattack occurs, you can quickly recover and restore data from Safeguarded copies.

Before you begin

If you are using IBM Storage Sentinel, you can scan snapshots for ransomware after you register a Security Scan Server by selecting Perform Security Scan. To register a Security Scan Server, see [Registering a Security Scan Server](#).

About this task

The Safeguarded Copy feature creates safeguarded snapshots that are not accessible by the host system and protects these snapshots from corruption that can occur in the production environment. You can define a Safeguarded Copy schedule to create multiple snapshots regularly, such as hourly or daily. You can also restore a snapshot to a different volume.

Procedure

1. Click the **Configure** tab. On the Views pane, select **Sites & Providers**, then expand the IBM Storage Virtualize for Snapshot and add the supported storage.
2. On the Views pane, select **SLA Policies**. The All SLA Policies pane opens.
3. In the All SLA Policies pane, click **New**. The New SLA Policies pane opens.
4. Add a subpolicy (SLA Policy) to an IBM Storage Virtualize for Snapshot SLA policy.
 - a) Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field, select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Note: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.
 - b) Click **Add Safeguarded Copy**.
 - c) In the Options pane, set the Safeguarded Copy subpolicy 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 subpolicy name displayed in IBM Storage Copy Data Management. The default name is Safeguarded Copy0.

Snapshot Prefix Label

Enter an optional label to identify the Snapshot. This label is added as a prefix to the Snapshot name created by the job.

Tip: Snapshot labels must contain only alphanumeric characters and underscores.

Perform Security Scan

You must enable **Perform Security Scan** and select your security scan servers. This allows you to scan for every specified backup number.

- d) Enter a name for the new subpolicy (SLA Policy).
- e) Click **Finish**.

Configure NetApp ONTAP SLA policies

The procedures describe how to configure NetApp ONTAP SLA policies.

Adding a snapshot sub-policy to a NetApp ONTAP SLA Policy

The procedures describe how to add a snapshot sub-policy to a NetApp ONTAP SLA Policy.

Procedure

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Tip: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add Snapshot**. In the Options pane, set the snapshot sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Disable system snapshot policy

Disables all the system snapshot jobs on the storage volumes.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Tip: 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 Storage Copy Data Management. The default initial name is Snapshot0.

Adding a VM Replication sub-policy to a NetApp ONTAP SLA Policy

The procedures describe how to add a VM Replication sub-policy to a NetApp ONTAP SLA Policy.

Procedure

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Tip: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add VM Copy**.
3. In the VM Replication Destination pane select an SVM from the list of available resources as the VM Replication destination, along with an associated aggregate. If no aggregate is selected, the aggregate with the largest amount of space available is chosen by default.
4. In the Options pane set the VM Replication sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Disable system snapshot policy

Disables all the system snapshot jobs on the storage volumes.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Tip: 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 Storage Copy Data Management. The default initial label is VM Copy0.

Storage Efficiency (Deduplication)

Enable or disable storage efficiency. Storage efficiency uses data deduplication to store the maximum amount of data while consuming less space.

Destination Datastore Type

Set the destination datastore type. Available datastore types include NFS and VMFS.

NFS - A NetApp volume will be created for NFS access and the target datastore will be created on that NFS share.

VMFS - A NetApp volume will be created and a LUN will be created on the volume. The volume will be mapped to the ESX, and the LUN will be formatted for VMFS. A VMFS datastore will be created on the LUN.

Adding a mirror sub-policy to a NetApp ONTAP SLA Policy

The procedures describe how to add a mirror sub-policy to a NetApp ONTAP SLA Policy.

Procedure

1. Select a snapshot, VM Replication, vault or mirror from the workflow pane and click **Add Mirror**.
2. In the Mirror Destination pane select a storage controller or SVM from the list of available resources as the mirror destination, along with an associated aggregate. If no aggregate is selected, the aggregate with the largest amount of space available is chosen by default.
3. In the Options pane set the **Mirror** sub-policy options.

Name

Enter an optional name to replace the default mirror sub-policy name displayed in IBM Storage 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 Storage Copy Data Management.

Volume Prefix Label

Enter an optional label to identify the volume. This label is added as a prefix to the volume name created by the job and cannot be edited after the job is submitted.

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

Storage Efficiency (Deduplication)

Enable or disable snapshot storage efficiency. Storage efficiency uses data deduplication to store the maximum amount of data while consuming less space.

Throttle

Set the transfer throughput in KBs per second between the source and the destination, which controls the number of parallel transfers that can take place.

Destination storage limit in GB / Destination volumes limit

Specify quotas for storage usage and the number of volume created on the destination for all jobs utilizing the SLA Policy.

Adding a vault sub-policy to a NetApp ONTAP SLA Policy

The procedures describe how to add a vault sub-policy to a NetApp ONTAP SLA Policy.

Procedure

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 Storage 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 Storage Copy Data Management.

Volume Prefix Label

Enter an optional label to identify the volume. This label is added as a prefix to the volume name created by the job and cannot be edited after the job is submitted.

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

Storage Efficiency (Deduplication)

Enable or disable snapshot storage efficiency. Storage efficiency uses data deduplication to store the maximum amount of data while consuming less space.

Throttle

Set the transfer throughput in KBs per second between the source and the destination, which controls the number of parallel transfers that can take place.

Destination storage limit in GB / Destination volumes limit

Specify quotas for storage usage and the number of volume created on the destination for all jobs utilizing the SLA Policy.

Configure Pure Storage FlashArray SLA policies

The procedures describe how to configure Pure Storage FlashArray SLA policies.

Adding a VM Replication sub-policy to a Pure Storage FlashArray SLA Policy

The procedures describe how to add a VM Replication sub-policy to a Pure Storage FlashArray SLA Policy.

Procedure

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Tip: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

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

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Target Volume Prefix Label

Enter an optional label to identify the target volume. This label is added as a prefix to the volume name created by the job.

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

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Tip: 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 Storage Copy Data Management. The default initial label is VM Copy0

Protocol

If more than one storage protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

Full Copy Method

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

Adding a snapshot sub-policy to a Pure Storage FlashArray SLA Policy

The procedures describe how to add a snapshot sub-policy to a Pure Storage FlashArray SLA Policy.

Procedure

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Tip: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

2. Click **Add Snapshot**. In the Options pane, set the snapshot sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Tip: 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 Storage Copy Data Management. The default initial name is Snapshot0.

3. If your Pure Storage FlashArray supports CloudSnap functionality, you can add snapshot offload which will create an offload copy on a S3 cloud storage target or NFS share. This option is only available after a snapshot sub-policy is added. Click **Add Snapshot Offload** or right-click on the snapshot sub-policy and click **Add Snapshot Offload**. In the Options pane, set the snapshot offload sub-policy options.

Keep Snapshots

After a certain number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots** field.

Tip: The IBM Storage Copy Data Management user interface will indicate that a Pure Storage FlashArray CloudSnap offload job has completed even though the transfer is still occurring in the background which is dependent on network speeds. Consider setting an age as the retention for offload copies when using the Pure Storage FlashArray CloudSnap functionality. Doing so will ensure that a sufficient amount of time has passed for data to be transferred to the S3 storage target or NFS share before it has to be condensed out from a backup. This is particularly important if several offload jobs are run in quick succession.

Offload Target

Select **Cloud** from the menu as the offload target.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Tip: 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 Storage Copy Data Management. The default initial name is Snapshotoffload0.

Adding a Replication sub-policy to a Pure Storage FlashArray SLA Policy

The procedures describe how to add a Replication sub-policy to a Pure Storage FlashArray SLA Policy.

Procedure

1. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, then set the interval in the **Interval** field. The lowest available frequency is five minutes.

Tip: Edits to the frequency and interval of an SLA Policy apply to all associated job schedules.

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

Keep Source Snapshots / Keep Destination Snapshots

A Pure Storage replication sub-policy provides snapshots to both a Source, or Primary location, and a Destination, or Replication location. After a certain number of snapshot instances are created for a resource, older instances are purged from the Source and Destination. In the **Keep Source Snapshots** and **Keep Destination Snapshots** fields, enter the age of the snapshot instances to purge in the **Days** field, or the number of instances to keep in the **Snapshots (maximum)** field.

Name

Enter an optional name to replace the default Replication sub-policy name displayed in IBM Storage Copy Data Management. The default initial name is Replication0.

Snapshot Prefix Label

Enter an optional label to identify the snapshot. This label is added as a prefix to the snapshot name created by the job.

Tip: Snapshot labels must contain only alphanumeric characters and underscores.

5. When you are satisfied that the SLA Policy-specific information is correct, click **Finish**. The SLA Policy appears on the All SLA Policies pane and can be applied to new and existing Backup job definitions.

Chapter 13. Configure scripts

Prescripts and postscripts are scripts that can be run before or after Backup and Restore jobs run, both at a job-level and before or after snapshots are captured. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

Scripts can be created locally, uploaded to your environment through the Scripts pane, then applied to job definitions. In a Windows environment, if your application supports VSS, the Backup job triggers the VSS application quiesce logic if the Make these VMs application/file system consistent option is enabled when creating the VMware Backup job. However, for applications that don't support VSS, or on Linux virtual machines, pre and post snapshot scripts can be used to quiesce your application for the snapshot backup.

Tip: If adding a script to a Windows-based File System job definition, the user running the script must have the "Log on as a service" right enabled, which is required for running prescripts and postscripts. For more information about the "Log on as a service" right, see [Add the Log on as a service Right to an Account](#).

Uploading a script

Supported scripts include shell scripts for Linux-based virtual machines and Batch and PowerShell scripts for Windows-based virtual machines. Scripts must be created using the operating system's associated file format.

Procedure

1. Click the **Configure** tab. On the **Views** panel, select **Scripts**.
2. Click **Upload**. The Upload Script dialog opens.
3. In the Script field, browse for a local script to upload, then click **Open**.
4. Enter an optional comment, then click **OK**. The script appears on the **Scripts** panel and can be applied to supported jobs.

Related concepts

["Job definition overview" on page 95](#)

You can create and edit job definitions from the **Jobs** tab. A job definition is a user-defined set of tasks and rules.

["Using state and status arguments in postscripts" on page 206](#)

Related reference

["Return code reference" on page 342](#)

Return Codes are issued when a script generated by a user-defined script is run. As the script runs, IBM Storage 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.

Replacing a script

Scripts can be created locally, uploaded to your environment through the Scripts pane, then applied to job definitions. Upload a revised version of a script.

Procedure

1. Click the **Configure** tab. On the **Views** panel, select **Scripts**.
2. In the Scripts panel, select the script to replace by clicking in the row containing the script name.

3. Click **Replace**. The Update Script dialog opens.
4. In the Script field, browse for a local updated script to upload, then click **Open**.
5. Enter an optional comment, then click **OK**. The revised script appears on the **Scripts** panel and can be applied to supported jobs.

Related concepts

[“Job definition overview” on page 95](#)

You can create and edit job definitions from the **Jobs** tab. A job definition is a user-defined set of tasks and rules.

[“Using state and status arguments in postscripts” on page 206](#)

Related reference

[“Return code reference” on page 342](#)

Return Codes are issued when a script generated by a user-defined script is run. As the script runs, IBM Storage 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.

Deleting a script

Delete a script when it becomes obsolete. Removing a script from an associated job definition allows you to delete the script immediately. Deleting the job definition while a script is still assigned to a job definition requires that you run the Maintenance job before deleting the script.

Procedure

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

Related concepts

[“Job definition overview” on page 95](#)

You can create and edit job definitions from the **Jobs** tab. A job definition is a user-defined set of tasks and rules.

[“Using state and status arguments in postscripts” on page 206](#)

Related reference

[“Return code reference” on page 342](#)

Return Codes are issued when a script generated by a user-defined script is run. As the script runs, IBM Storage 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.

Chapter 14. Schedules

The topics in the following section cover creating, editing, and deleting schedules.

Creating a schedule

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

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](#).

About this task

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

Tip: 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.

Procedure

1. Click the **Configure** tab. On the Views pane, select **Schedules**.
2. In the Schedules pane, click **New**. The Properties pane displays.
3. Revise fields on the Properties pane:
 - a) In **Name**, enter a descriptive schedule name. By default, your schedule parameters are added to the schedule description.
 - b) In **Trigger**, select the frequency for the job session to run:
 - Once** to schedule a job session to run once. In **Trigger date**, select the day of the week for the job session to run. In **Time of day**, select a starting time.
 - Hourly** to schedule a job session to run hourly. In **Interval**, select the number of hours between job sessions. In **Time of day** and **Starts**, select a starting time and date. If applicable, enter an expiration date in **Expires**.
 - Daily** to schedule a job session to run daily or every few days. In **Interval**, select the number of days between job sessions. In **Time of day** and **Starts**, select a starting time and date. If applicable, enter an expiration date in **Expires**.
 - Weekly** to schedule a job session to run weekly or every few weeks. In **Day of week**, select the day of the week for the job session to run during the week. In **Time of day** and **Starts**, select a starting time and date. If applicable, enter an expiration date in **Expires**.
 - Monthly** to schedule a job session to run monthly or every few months. In **Day of month**, select the day or days of the month for the job session to run during the month. In **Time of day** and **Starts**, select a starting time and date. If applicable, enter an expiration date in **Expires**.
4. Click **Save**. The revisions are applied to the schedule.

What to do next

Assign the schedule to a new or existing job definition. See [Jobs Overview](#) and [Edit a Job Definition](#).

Related tasks

[“Editing a schedule” on page 88](#)

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.

[“Deleting a schedule” on page 89](#)

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

Editing 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.

Before you begin

Review the properties of your current schedules. See [Create a Schedule](#).

About this task

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.

Tip: 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.

Procedure

1. Click the **Configure** tab. On the Views pane, select **Schedules**.
2. Select the schedule to edit by clicking in the row containing the schedule name.
3. Revise fields on the Properties pane:
 - a) In **Name**, enter a descriptive schedule name. By default, your schedule parameters are added to the schedule description.
 - b) In **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 revisions are applied to the schedule.

What to do next

Assign the edited schedule to a new or existing job definition. See [Jobs Overview](#) on page 165 and [Edit a Job Definition](#).

Related tasks

[“Creating a schedule” on page 87](#)

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

[“Deleting a schedule” on page 89](#)

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

Deleting a schedule

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

About this task

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

Procedure

1. Click the **Configure** tab. On the Views pane, select **Schedules**.
2. Select the schedule to delete by clicking in the row containing the schedule name.
3. Click **Delete**. A confirmation dialog box opens.
4. Confirm deletion. The schedule is deleted.

Related tasks

[“Creating a schedule” on page 87](#)

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

[“Editing a schedule” on page 88](#)

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.

Chapter 15. Jobs

The topics in the following section cover defining, editing, and deleting job definitions, as well as descriptions of job types.

Jobs overview

You can create and edit job definitions from the **Jobs** tab. You can also start, monitor, stop, and resume job sessions from the Jobs tab. From this pane, you can also view all scheduled and unscheduled job sessions, and start jobs before scheduled run times.

A job definition is a user-defined set of tasks and rules. Once a job definition is added to IBM Storage Copy Data Management, it can be combined with a schedule or trigger to create a job. There are several job types including Inventory, Backup, Restore, Reports, and Scripts

Remember: If the IBM Storage Copy Data Management virtual appliance is shut down while jobs are in progress, the jobs will automatically restart once the virtual appliance is back online. By default, jobs that began running 30 minutes before the appliance restarted will automatically restart.

In addition, you can view the job details. Select a job to view the current job status, the job schedule, and control the activity for the selected job.

Select a job and click **View Last Run Session** to view the duration and completion status of the most recent run of a selected job.

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

Related concepts

[“Job definition overview” on page 95](#)

You can create and edit job definitions from the **Jobs** tab. A job definition is a user-defined set of tasks and rules.

[“Monitor a job session” on page 93](#)

You can view the details of a job session that is currently running or one that has finished.

[“Start, pause, and hold a job session” on page 91](#)

From the Jobs tab you can run a job session on demand, pause or cancel a running job, and hold all future scheduled instances of a job from running until you are ready for the job to proceed.

Start, pause, and hold a job session

From the Jobs tab you can run a job session on demand, pause or cancel a running job, and hold all future scheduled instances of a job from running until you are ready for the job to proceed.

If the IBM Storage Copy Data Management virtual appliance is shut down while jobs are in progress, the jobs will automatically restart once the virtual appliance is back online. By default, jobs that began running 30 minutes before the appliance restarted will automatically restart.

Starting a job session

The procedures describe how to start a job session.

Procedure

1. Click the **Jobs** tab.
2. Select the job to run by clicking in the row containing the job name.
3. Click **Start**, or right-click the job name and select **Start**. A confirmation dialog box opens.

Tip: If a job session has multiple run options, such as running a job session in Test, Recovery, or Clone mode, you will be prompted to select a job session type.

4. Click **Yes**. The job session runs.
5. In the **Activity** pane, click the job name to view the job session details, including the job session's start date and time, duration, description, status through a progress bar, and associated messages.

What to do next

- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#).
- Once the job session completes, review cataloged data through the Search and Report tabs. See [Browse Inventory](#) and [Report Overview](#).

Related tasks

[“Creating a schedule” on page 87](#)

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

Pausing and resuming a job session

The procedures describe how to pause and resume a job session.

Procedure

1. Click the **Jobs** tab.
2. Select a running job to stop by clicking in the row containing the job name.
3. From the **More Actions** drop-down menu, select **Pause**, or right-click the job name and select **Pause**. A confirmation dialog box opens.
4. Click **Yes**. The job session pauses.
5. Select **Resume** from the **More Actions** drop-down menu, or right-click the job name and select **Resume** to resume the job session.

What to do next

- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#).
- Once the job session completes, review cataloged data through the Search and Report tabs. See [Browse Inventory](#) and [Report Overview](#).

Related tasks

[“Creating a schedule” on page 87](#)

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

Holding and releasing a job session

The procedures describe how to hold and release a job session.

Procedure

1. Click the **Jobs** tab.
2. Select the job to suspend by clicking in the row containing the job name.
3. From the **More Actions** drop-down menu, select **Hold Schedule**, or right-click the job name and select **Hold Schedule** to hold the job. The job session status changes to **Held**, and all future scheduled instances of the job will not run until released.
4. From the **More Actions** drop-down menu, select **Release Schedule**, or right-click the job name and select **Release Schedule** to release the job.

What to do next

- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#).
- Once the job session completes, review cataloged data through the Search and Report tabs. See [Browse Inventory](#) and [Report Overview](#).

Related tasks

[“Creating a schedule” on page 87](#)

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

Canceling a job session

The procedures describe how to cancel a job session.

Procedure

1. Click the **Jobs** tab.
2. Select the job session to cancel by clicking in the row containing the job name.
3. From the **More Actions** drop-down menu, select **Cancel** to cancel the job. The job session status changes to **Canceled**.

What to do next

- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#).
- Once the job session completes, review cataloged data through the Search and Report tabs. See [Browse Inventory](#) and [Report Overview](#).

Related tasks

[“Creating a schedule” on page 87](#)

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

Monitor a job session

You can view the details of a job session that is currently running or one that has finished.

Use the Jobs tab to view the status of a job session including start time, end time, and job name.

Monitoring a running job session

The procedures describe how to monitor a running job session.

Before you begin

Start a job session. If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#).

Procedure

1. Click the **Jobs** tab. The All Jobs folder displays defined jobs that are currently running or idle, and provides information about their most recent session, their last runtime, last run duration, and last run status. To view a specific job type, select the associated job type from the folder structure.
2. View the status of a job in the status column. Currently running job sessions are represented by an active icon. Once a job session finishes, one of the following icons appears in the status column:
 - a) **Completed** - Indicates the job session completed successfully. All tasks associated with the job session were completed
 - b) **Partial** - Indicates the job session completed, but one or more tasks failed or were skipped.

- c) **Failed** - Indicates the job session did not successfully complete due to mixed task statuses.
- d) **Aborted** - Indicates the job session did not successfully complete due to a reset, reboot, or shutdown of the virtual appliance server.
- e) **Held** - Indicates the job has been paused through the Halt feature in the Actions menu.
- f) **Idle** - Indicates the job session is idle.
- g) **Skipped** - Indicates that a volume was not cataloged. See the Task tab for more information about skipped jobs.
- h) **Stopped** - Indicates the job was stopped using the Stop button.

Related concepts

[“Start, pause, and hold a job session” on page 91](#)

From the Jobs tab you can run a job session on demand, pause or cancel a running job, and hold all future scheduled instances of a job from running until you are ready for the job to proceed.

[“Collecting logs for troubleshooting” on page 317](#)

For troubleshooting the IBM Storage Copy Data Management application, IBM Storage Copy Data Management can generate an archive of logs containing various files.

Filtering the list of jobs based on type or status

The procedures describe how to filter the list of jobs based on type or status.

Before you begin

Start a job session. If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#).

Procedure

1. Click the **Jobs** tab.
2. Click the drop-down arrow in the header of the **Status** or **Last Run Status** columns.
3. Select **Filters**, then choose a filter criteria.

Related concepts

[“Start, pause, and hold a job session” on page 91](#)

From the Jobs tab you can run a job session on demand, pause or cancel a running job, and hold all future scheduled instances of a job from running until you are ready for the job to proceed.

[“Collecting logs for troubleshooting” on page 317](#)

For troubleshooting the IBM Storage Copy Data Management application, IBM Storage Copy Data Management can generate an archive of logs containing various files.

Viewing information about specific job sessions

The procedures describe how to view information about specific job sessions.

Before you begin

Start a job session. If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#).

Procedure

1. Click the **Jobs** tab and select a job by clicking in the row containing the job name.
2. The Activity / History pane displays. If the job is currently running, its status displays in the **Activity** tab. To review the history of previous runs of the job, click the **History** tab, then click the session for details. The following panes open:

Tasks

Displays a task-by-task view of the job session, including start and end times, duration, and status. It also displays details of the underlying tasks that take place during the job session, including the task's type, duration, and status.

Details

Displays an overview of the job definition, including the job name, sources, options, and notification settings.

Log

Displays the job log, which can be used for troubleshooting purposes.

Related concepts

[“Start, pause, and hold a job session” on page 91](#)

From the Jobs tab you can run a job session on demand, pause or cancel a running job, and hold all future scheduled instances of a job from running until you are ready for the job to proceed.

[“Collecting logs for troubleshooting” on page 317](#)

For troubleshooting the IBM Storage Copy Data Management application, IBM Storage Copy Data Management can generate an archive of logs containing various files.

Job definition overview

You can create and edit job definitions from the **Jobs** tab. A job definition is a user-defined set of tasks and rules.

Once a job definition is added to IBM Storage Copy Data Management, it can be combined with a schedule or trigger to create a job. There are several job types including Inventory, Backup, Restore, Reports, and Scripts.

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

When a job is run, the job session status is monitored and its status can be watched real-time in the job monitor. Job sessions are run on demand or by trigger.

Tip: Create a schedule before creating a job definition so that you can easily add the schedule to the job definition.

Job types

Inventory jobs

Inventory jobs interrogate storage systems to gather and record metadata about high-level objects and files. You can select one or more providers of the same type in a single job definition for cataloging.

Backup and Restore jobs

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

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

Report jobs

A Report job is a System job that summarizes information about cataloged providers and the data and other resources that reside on them.

Script jobs

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

Maintenance job

The Maintenance job removes resources and associated objects created by IBM Storage Copy Data Management when a job in a pending state is deleted. The cleanup procedure reclaims space on your storage devices, cleans up your IBM Storage Copy Data Management catalog, and removes related snapshots.

Related concepts

[“Configure SLA policies” on page 73](#)

[“Monitor a job session” on page 93](#)

You can view the details of a job session that is currently running or one that has finished.

[“Creating an IBM Storage Virtualize Restore job definition” on page 169](#)

[“Creating an NetAPP ONTAP Restore job definition” on page 179](#)

[“Creating a Pure Storage FlashArray Restore job definition” on page 187](#)

[“Creating a VMware Restore job definition” on page 191](#)

Related tasks

[“Creating a schedule” on page 87](#)

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

[“Creating a Database Inventory job definition” on page 97](#)

A Database Inventory job provides the framework to catalog and collect information about your application database servers.

[“Creating a File System Inventory job definition” on page 99](#)

A File System Inventory job provides the framework to catalog and collect information about your physical Windows, Linux and AIX file systems.

[“Creating an IBM Storage Virtualize Inventory job definition” on page 100](#)

An IBM Storage Virtualize Inventory job provides the framework to catalog and collect information about high-level IBM objects on your IBM storage systems.

[“Creating a NetApp ONTAP Storage Inventory job definition” on page 103](#)

A NetApp ONTAP Storage Inventory job provides the framework to catalog and collect information about high-level NetApp ONTAP objects.

[“Creating a NetApp ONTAP File Inventory job definition” on page 104](#)

A NetApp ONTAP File Inventory job provides the framework to catalog and collect information about low-level NetApp ONTAP objects.

[“Creating a Pure Storage FlashArray Inventory job definition” on page 107](#)

A Pure Storage FlashArray Inventory job provides the framework to catalog and collect information about high-level objects on your Pure Storage systems.

[“Creating a VMware Inventory job definition” on page 109](#)

A VMware Inventory job provides the framework to catalog and collect information about VMware objects.

[“Creating an InterSystems Database Backup job definition” on page 110](#)

The procedures describe how to create an InterSystems Database Backup job definition. You can create InterSystems Caché and InterSystems IRIS application servers in IBM Storage Copy Data Management as an InterSystems Database.

[“Creating a SAP HANA Backup job definition” on page 114](#)

This procedure describes how to create a SAP HANA Backup job definition.

[“Creating an Oracle Backup job definition” on page 116](#)

The procedures describe how to create an Oracle Backup job definition.

[“Creating a SQL Backup job definition” on page 125](#)

The procedures describe how to create a SQL Backup job definition.

[“Creating a File System Backup job” on page 129](#)

The procedures describe how to create a File System Backup job definition.

[“Creating an IBM Storage Virtualize Backup job definition” on page 131](#)

The procedures describe how to create an IBM Storage Virtualize Backup job definition.

[“Creating a NetApp ONTAP Backup job definition” on page 140](#)

The procedures describe how to create a NetApp ONTAP Backup job definition.

[“Creating a Pure Storage FlashArray Backup job definition” on page 142](#)

The procedures describe how to create a Pure Storage FlashArray Backup job definition.

[“Creating a VMware Backup job definition” on page 145](#)

The procedures describe how to create a VMware Backup job definition.

[“Creating an InterSystems Database Restore job definition” on page 151](#)

The procedures describe how to create an InterSystems Database Restore job definition. You can create InterSystems Caché and InterSystems IRIS application servers in IBM Storage Copy Data Management as an InterSystems Database.

[“Creating a SAP HANA Restore job definition” on page 154](#)

This procedure describes how to create a SAP HANA Restore job definition.

[“Creating an Oracle Restore job definition” on page 157](#)

The procedures describe how to create an Oracle Restore job definition.

[“Creating a Microsoft SQL Restore job definition” on page 162](#)

The procedures describe how to create a Microsoft SQL Restore job definition.

[“Creating a File System Restore job definition” on page 166](#)

The procedures describe how to create a File System Restore job definition.

Inventory jobs

The topics in the following section cover Inventory job definitions.

Creating a Database Inventory job definition

A Database Inventory job provides the framework to catalog and collect information about your application database servers.

Before you begin

- At least one application provider must be associated with a Database Inventory job definition, such as an Oracle resource. Before defining an Inventory job, add application providers. See [Register a Provider](#).
- Credentials are required for cataloging application servers. See [Identities Overview](#).
- Review database requirements. See [Oracle requirements](#), [Microsoft SQL requirements](#), [SAP HANA requirements](#), [InterSystems database requirements](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

- If an Oracle Inventory job runs at the same time or short period after an Oracle Backup job runs, copy errors may occur due to temporary mounts that are created during the Backup job. As a best practice, schedule Oracle Inventory jobs so that they do not overlap with Oracle Backup jobs.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#).

About this task

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

A Database Inventory job provides the framework to catalog and collect information about your application database servers.

Run a Database Inventory job to analyze your application servers in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer sprawl, storage overutilization, and other storage inefficiencies.

Tip: To successfully catalog virtual providers, you must also register associated vCenter providers.

Procedure

1. Click the **Jobs** tab. Expand the **Database** folder, then select an application database type to catalog. Available application database types include **InterSystems Caché** and **InterSystems IRIS** (collectively referred to as **InterSystems Database**), **Oracle**, **SAP HANA**, and **Microsoft SQL**.
2. Click **New**, then select **Inventory**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more providers containing application data that you wish to catalog.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

Number of catalog instances to keep

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

7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

8. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

9. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

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

What to do next

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

Related tasks

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Creating a File System Inventory job definition

A File System Inventory job provides the framework to catalog and collect information about your physical Windows, Linux and AIX file systems.

Before you begin

- Review File System requirements. See [File system requirements](#).
- At least one file system provider must be associated with a File System Inventory job definition. Before defining an Inventory job, add file system providers. See [Register a Provider](#).
- Credentials are required for cataloging application servers. See [Identities Overview](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

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

About this task

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

A File System Inventory job provides the framework to catalog and collect information about your physical Windows, Linux and AIX file systems. You can select one or more file system providers of the same type in a single job definition for cataloging.

Run a File System Inventory job to analyze your file systems in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer sprawl, storage overutilization, and other storage inefficiencies.

Procedure

1. Click the **Jobs** tab. Expand the **File System** folder, then select **File System**.
2. Click **New**, then select **Inventory**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more file systems that you wish to catalog.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the Jobs tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

Number of catalog instances to keep

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

7. Optionally, expand the Notification section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click Add to add it to the list.

- Optionally, expand the **Schedule** section to select the job scheduling options. Select Start job now to create a job definition that starts the job immediately. Select Schedule job to start at later time to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

- When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

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

What to do next

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

Related tasks

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Creating an IBM Storage Virtualize Inventory job definition

An IBM Storage Virtualize Inventory job provides the framework to catalog and collect information about high-level IBM objects on your IBM storage systems.

Before you begin

- At least one application provider must be associated with a Database Inventory job definition, such as an Oracle resource. Before defining an Inventory job, add application providers. See [Register a Provider](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

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

About this task

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

An IBM Storage Virtualize Inventory job provides the framework to catalog and collect information about high-level IBM objects on your IBM storage systems. You can select one or more IBM providers in a single job definition for cataloging.

Run an IBM Storage Virtualize Inventory job to analyze your IBM environment in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer sprawl, storage overutilization, and other storage inefficiencies.

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **IBM Storage Virtualize**.
2. Click **New**, then select **Inventory**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more providers to catalog.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

Number of catalog instances to keep

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

7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

8. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

9. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

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

What to do next

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

Related tasks

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Creating an IBM Storage Virtualize for Snapshot Inventory job definition

An IBM Storage Virtualize for Snapshot Inventory job provides the framework to catalog and collect information about high-level IBM objects on your IBM storage systems.

Before you begin

- At least one application provider must be associated with a Database Inventory job definition, such as an Oracle resource. Before defining an Inventory job, add application providers. See [Register a Provider](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

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

About this task

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

An IBM Storage Virtualize for Snapshot Inventory job provides the framework to catalog and collect information about high-level IBM objects on your IBM storage systems. You can select one or more IBM providers in a single job definition for cataloging.

Run an IBM Storage Virtualize for Snapshot Inventory job to analyze your IBM environment in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer sprawl, storage overutilization, and other storage inefficiencies.

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **IBM Storage Virtualize for Snapshot**.
2. Click **New**, then select **Inventory**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more providers to catalog.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

Number of catalog instances to keep

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

7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

8. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

9. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

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

What to do next

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

Creating a NetApp ONTAP Storage Inventory job definition

A NetApp ONTAP Storage Inventory job provides the framework to catalog and collect information about high-level NetApp ONTAP objects.

Before you begin

- At least one application provider must be associated with a Database Inventory job definition, such as an Oracle resource. Before defining an Inventory job, add application providers. See [Register a Provider](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

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

About this task

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

A NetApp ONTAP Storage Inventory job provides the framework to catalog and collect information about high-level NetApp ONTAP objects. You can select one or more NetApp ONTAP cluster or non-cluster providers in a single job definition for cataloging. NetApp ONTAP 7-Mode and Cluster-Mode are both supported.

Run a NetApp ONTAP Storage Inventory job to analyze your NetApp ONTAP environment in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer snapshot sprawl, storage overutilization, and other storage inefficiencies.

Tip: For cataloging low-level NetApp ONTAP objects, create a NetApp ONTAP File Inventory job definition.

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **NetApp ONTAP**.
2. Click **New**, then select **NetApp ONTAP Storage**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more providers to catalog.

5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

Connection timeout (secs)

To run a catalog job, the application needs to connect with the resource. If there is no response within a certain time limit, it times out and the job session fails. Enter the number of seconds to wait before timing out.

Number of catalog instances to keep

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

7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

8. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

9. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

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

What to do next

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

Related tasks

[“Creating a NetApp ONTAP File Inventory job definition” on page 104](#)

A NetApp ONTAP File Inventory job provides the framework to catalog and collect information about low-level NetApp ONTAP objects.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

Creating a NetApp ONTAP File Inventory job definition

A NetApp ONTAP File Inventory job provides the framework to catalog and collect information about low-level NetApp ONTAP objects.

Before you begin

- At least one application provider must be associated with a Database Inventory job definition, such as an Oracle resource. Before defining an Inventory job, add application providers. See [Register a Provider](#).

- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

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

About this task

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

A NetApp ONTAP File Inventory job provides the framework to catalog and collect information about low-level NetApp ONTAP objects. You can select one or more NetApp ONTAP providers in a single job definition for cataloging. NetApp ONTAP 7-Mode and Cluster-Mode are both supported.

During cataloging, properties for NetApp ONTAP files are collected and stored. Run a NetApp ONTAP File Inventory job to analyze your low-level NetApp ONTAP environment in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer snapshot sprawl, storage overutilization, and other storage inefficiencies.

Tip:

- For cataloging low-level NetApp ONTAP objects, create a NetApp ONTAP File Inventory job definition.
- NetApp ONTAP File Inventory jobs may take a long time to run because they catalog at a file level. Consider creating NetApp ONTAP File Inventory job definitions that constrain the number of file system objects to catalog by limiting the number of storage systems and volume processed by a single job. You can also stagger the scheduled run times of the jobs to prevent them from running concurrently. Consider running NetApp ONTAP File Inventory jobs less frequently than you run NetApp ONTAP Storage Inventory jobs.
- Create a schedule before creating a job definition so that you can easily add the schedule to the job definition.

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **NetApp ONTAP**.
2. Click **New**, then select **NetApp ONTAP Files**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more volumes to catalog. For Cluster-Mode providers, the SVM name appears in parentheses after the volume name. To view the number of files on the selected volume, hover your cursor over the volume name.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

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.

Skip root volume

Select this option to avoid cataloging the root volume of your NetApp ONTAP objects.

Skip unsupported volumes

Select this option to avoid cataloging unsupported volumes such as volumes that are offline. Selecting this option also avoids cataloging volumes with i2p disabled when Catalog all available snapshots is set to Yes and Traversal Method is set to Snapdiff.

Connection timeout (secs)

To run a catalog job, the application needs to connect with the resource. If there is no response within a certain time limit, it times out and the job session fails. Enter the number of seconds to wait before timing out.

Number of catalog instances to keep

After a certain number of job runs for a given job, older low-level NetApp ONTAP objects for that job are purged from the Inventory. Enter the number of job runs for which low-level NetApp ONTAP objects are to be retained. Note that by default, the old data for the job is purged from the IBM Storage Copy Data Management Inventory after the newer data is cataloged.

Condense catalog before run

Select this option to purge older low-level NetApp ONTAP objects from the IBM Storage Copy Data Management Inventory before newer data is cataloged through a NetApp ONTAP File Inventory job.

Condense catalog after failed run

Select this option to purge NetApp ONTAP objects from the IBM Storage Copy Data Management Inventory for the unsuccessful run of the NetApp ONTAP File Inventory job.

Traversal Method

This option indicates the methodology to employ when cataloging snapshots. IBM Storage 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.

Note: SnapDiff as Traversal Method is no longer supported.

Filewalk

IBM Storage 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 Storage Copy Data Management Filewalker tool. Once cataloging completes, searchable owner information is added to the Inventory. See the Filewalker documentation for more information.

Base snapshot for catalog

For file cataloging, a base snapshot on the storage system is required.

Create a new snapshot before cataloging. A new snapshot is created when the job session begins and is deleted after cataloging. This is the preferred method if sufficient resources are available.

Use latest snapshot. A new snapshot is not created; instead the latest healthy snapshot available for that volume is used to assist with file level cataloging. You can exclude certain snapshots from the set of snapshots IBM Storage Copy Data Management selects from by entering character patterns for those snapshots. Use a comma to separate each pattern.

Tip: To catalog a volume that does not support snapshot creation, such as a SnapMirror volume, select Use latest snapshot.

Catalog all available snapshots

Yes. Catalogs snapshots on the volume in addition to the base snapshot. IBM Storage Copy Data Management examines additional snapshots ensuring the Inventory has the latest data. You can exclude certain snapshots from the set of snapshots IBM Storage Copy Data Management catalogs by entering character patterns for those snapshots. Use a comma to separate each pattern.

No. Catalogs only the base snapshot.

Tip: By cataloging all available snapshots, you can view multiple versions of your files through a file-level search. Additional versions are available on the file's properties pane.

7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

8. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

9. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

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

What to do next

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

Related tasks

[“Creating a NetApp ONTAP Storage Inventory job definition” on page 103](#)

A NetApp ONTAP Storage Inventory job provides the framework to catalog and collect information about high-level NetApp ONTAP objects.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Related reference

[“Search and filter guidelines” on page 339](#)

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.

Creating a Pure Storage FlashArray Inventory job definition

A Pure Storage FlashArray Inventory job provides the framework to catalog and collect information about high-level objects on your Pure Storage systems.

Before you begin

- At least one Pure Storage provider must be associated with a Pure Storage FlashArray Inventory job definition. Before defining an Inventory job, add Pure Storage providers. See [Register a Provider](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

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

About this task

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

A Pure Storage FlashArray Inventory job provides the framework to catalog and collect information about high-level objects on your Pure Storage systems. You can select one or more Pure Storage providers in a single job definition for cataloging.

Run a Pure Storage FlashArray Inventory job to analyze your Pure Storage environment in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer sprawl, storage overutilization, and other storage inefficiencies.

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **Pure Storage FlashArray**.
2. Click **New**, then select **Inventory**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more providers containing application data that you wish to catalog.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

Number of catalog instances to keep

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

7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

8. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

9. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

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

What to do next

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

Related tasks

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Creating a VMware Inventory job definition

A VMware Inventory job provides the framework to catalog and collect information about VMware objects.

Before you begin

- At least one VMware provider must be associated with a VMware Inventory job definition. Before defining an Inventory job, add VMware providers. See [Register a Provider](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

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

About this task

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

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

Run a VMware Inventory job to analyze your VMware environment in real time and navigate and correlate the objects from across the Enterprise in a single view. Additionally, you can infer snapshot sprawl, storage overutilization, and other storage inefficiencies.

Procedure

1. Click the **Jobs** tab. Expand the **Hypervisor** folder, then select **VMware**.
2. Click **New**, then select **Inventory**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites, select one or more providers containing application data that you wish to catalog.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
6. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that can be performed on the provider.

Connection timeout (secs)

To run a catalog job, the application needs to connect with the resource. If there is no response within a certain time limit, it times out and the job session fails. Enter the number of seconds to wait before timing out.

Number of catalog instances to keep

After a certain number of job runs for a given job, older VMware objects for that job are purged from the Inventory. Enter the number of job runs for which VMware objects are to be retained.

7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

8. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

9. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

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

What to do next

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

Related tasks

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Backup jobs

The topics in the following section cover Backup job definitions as well as creating VMware Backup job proxies.

Creating an InterSystems Database Backup job definition

The procedures describe how to create an InterSystems Database Backup job definition. You can create InterSystems Caché and InterSystems IRIS application servers in IBM Storage Copy Data Management as an InterSystems Database.

Before you begin

- Create and run a Database Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - Database](#).
- Configure an SLA Policy. See [Configure SLA Policies](#).
- Review InterSystems Database requirements. See [InterSystems database requirements](#).

Considerations:

- InterSystems Database backup jobs occur at the instance level.
- Note that it is possible to scan in an InterSystems Database backup failover member instance or an async member instance and run snapshots against the mirror copy instead of the primary failover member.

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

Important: You should not run Incremental FC and FC NoCopy backup jobs on the same source volume.

About this task

IBM Storage Copy Data Management provides application database copy management through application-consistent backup creation, cloning, and recovery. IBM Storage Copy Data Management copy management leverages the snapshot and replication features of the underlying storage platform to create, replicate, clone, and restore backups of InterSystems Database instances.

IBM Storage Copy Data Management auto-discovers databases and enables backups only of eligible instances. To be eligible for backup, instances must reside on supported storage platforms.

Procedure

1. Click the **Jobs** tab. Expand the **Database** folder, then select **InterSystems Database**.
2. Click **New**, then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more resources to back up.
5. Select an SLA Policy that meets your backup data criteria.

Tip: For Sentinel scanning capabilities, the SLA policy must be Safeguarded Copy and you need to select a previously registered Security Scan server.

6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs** tab. Repeat as necessary to add additional SLA Policies to the job definition. If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Tip: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Skip IA[®] Mount points and/or databases

Enable to skip Instant Disk Restore objects. By default, this option is enabled.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#) on page 156.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

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

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

10. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.

What to do next

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

Related tasks

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

[“Creating an InterSystems Database Restore job definition” on page 151](#)

The procedures describe how to create an InterSystems Database Restore job definition. You can create InterSystems Caché and InterSystems IRIS application servers in IBM Storage Copy Data Management as an InterSystems Database.

Creating an InterSystem IRIS application (on AIX) scanning backup job

The procedures describe how to create an AIX-based application (InterSystem IRIS) Backup job definition for scanning for Sentinel.

Before you begin

- Create and run a Database Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - Database](#).
- Configure an SLA Policy. See [Configure SLA Policies](#).
- Register at least one AIX proxy host as file system in the application. When the user selects **Enable Proxy Server** all the AIX proxy host(s) that are registered in the IBM Storage Copy Data Management are displayed.
- Configure the NFS version 3 or 4 (preferably 4) on the AIX proxy and Security Scan server.

Note: If NFS mounting fails during the backup job execution, perform the cleanups manually.

- A system property `protection.application.scan.fullNFSRestriction` has been introduced in `/opt/virgo/repository/ecx-usr/com.syncsort.dp.xsb.serviceprovider.properties`. By default the system property is set to

true to create secure NFS for the Security Scan server, but if it is set to **false** IBM Storage Copy Data Management will create a global NFS share.

- Make sure that the proxy and scan host can ping each other's host addresses which are configured in IBM Storage Copy Data Management.
- Make sure that the proxy host is zoned and added to the SVC array.

Note: AIX proxy is only utilized for Sentinel workloads.

Procedure

1. Click the **Jobs** tab. Expand the **Database** folder, then select **InterSystems Database**.
2. Click **New**, then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more resources to back up.
5. Select an SLA Policy that meets your backup data criteria.

Tip: SLA policy must be Safeguarded Copy and previously registered Security Scan server needs to be selected.

6. Select **Enable Proxy Server** check box. All the AIX proxy host(s) that are registered in the IBM Storage Copy Data Management are displayed.
7. Choose the appropriate proxy host(s) for IBM Storage Sentinel Anomaly Scan software.
8. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs** tab. Repeat as necessary to add additional SLA Policies to the job definition. If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Tip: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

9. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.
10. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Skip IA Mount points and/or databases

Enable to skip Instant Disk Restore objects. By default, this option is enabled.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#) on page 156.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

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

11. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.

What to do next

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Hold, and Cancel a Job Session](#).
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#).
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Create an AIX-based Application Restore job definition. See [Create a Restore Job Definition - InterSystems Database](#).

Creating a SAP HANA Backup job definition

This procedure describes how to create a SAP HANA Backup job definition.

Before you begin

- Review SAP HANA prerequisites. See [SAP HANA requirements](#).
- Create and run a Database Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - Database](#).
- Configure an SLA Policy. See [Configure SLA Policies](#).

Considerations:

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

Important: You should not run Incremental FlashCopy and FlashCopy No Copy backup jobs on the same source volume.

About this task

IBM Storage Copy Data Management provides application database copy management through application-consistent backup creation, and recovery. IBM Storage Copy Data Management leverages the snapshot and replication features of the underlying storage platform to create, replicate, and restore backups of SAP HANA databases.

IBM Storage Copy Data Management auto-discovers databases and enables backups only of eligible databases. To be eligible for backup, application databases must reside on supported storage platforms.

The log backup feature enables continuous backup of Archive logs to a specified destination. IBM Storage Copy Data Management leverages archived logs to enable point-in-time recoveries of databases to facilitate RPOs.

Procedure

1. Click the **Jobs** tab. Expand the **Database** folder, then select **SAP HANA**.
2. Click **New**, then select **Backup**. The job editor opens.

3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more resources to back up.

Tip: You cannot select a database if it is not eligible for protection. Hover your cursor over the database name to view the reasons the database is ineligible, such as the database files, control files, or redo log files are stored on unsupported storage.

5. Select an SLA Policy that meets your backup data criteria.

Tip: For Sentinel scanning capabilities, the SLA policy must be Safeguarded Copy and you need to select a previously registered Security Scan server.

6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs** tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Tip: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Skip IA Mount points and/or databases

Enable to skip Instant Disk Restore objects. By default, this option is enabled.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Scripts can also be uploaded and edited through the Scripts view on the **Configure** tab. For more information, see [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

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

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

10. To edit Log Backup options before creating the job definition, click Log Backup. If Backup Logs is selected, IBM Storage Copy Data Management backs up database logs then protects the underlying disks. Select resources in the Select resource(s) to add archive log destination field. Database logs are backed up to the directory entered in the Use Universal Destination Mount Point field, or in the Mount Point field after resources are selected. The destination must already exist, must reside on storage from a supported vendor, and the SAP HANA user needs to have full read and write access. For more information, see the [SAP HANA requirements](#).

If multiple databases are selected for backup, then each of the servers hosting the databases must have their Destination Mount Points set individually. For example, if two databases, one from Server A and one from Server B, are added to the same job definition, and a single mount point named /logbackup is defined in the job definition, then you must create separate disks for each server and mount them both to /logbackup on the individual servers. When the mount point is changed, you must manually go in and clean up the previous log backup directory path.

To disable a log backup schedule on the SAP HANA server, edit the associated SAP HANA Backup job definition and deselect the checkbox next to the database on which you wish to disable the log backup schedule in the Select resource(s) for log backup destination field, then save and re-run the job. When the mount point is disabled, you must manually go in and clean up the log backup directory path.

Tip:

The job definition must be saved and re-run for mount point changes or disablement to take effect.

The default setting for pruning SAP HANA log backups is 7 days.

This value may be adjusted in the property file located in /opt/virgo/repository/ecx-usr/com.syncsort.dp.xsb.serviceprovider.properties. Modify the application.logpurge.days parameter to the desired value. Finally, restart the virgo service by issuing the following command:

```
systemctl restart virgo.service
```

11. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the Jobs tab.

What to do next

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

Related tasks

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

[“Creating a SAP HANA Restore job definition” on page 154](#)

This procedure describes how to create a SAP HANA Restore job definition.

Creating an Oracle Backup job definition

The procedures describe how to create an Oracle Backup job definition.

Before you begin

- Create and run a Database Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - Database](#).
- Configure an SLA Policy. See [Configure SLA Policies](#).
- Review Oracle requirements. See [Oracle requirements](#).

Oracle database considerations:

- To ensure that filesystem permissions are retained correctly when IBM Storage Copy Data Management moves Oracle data between servers, ensure that the user and group IDs of the Oracle users (e.g. oracle, oinstall, dba) are consistent across all the servers. Refer to Oracle documentation for recommended uid and gid values.
- If Oracle data resides on LVM volumes, you must stop and disable the lvm2-lvmetad service before running Backup or Restore jobs. Leaving the service enabled can prevent volume groups from being resigntured correctly during restore and can lead to data corruption if the original volume group is also present on the same system. To disable the lvm2-lvmetad service, run the following commands:

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

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

```
use_lvmetad = 0
```

- Note that Oracle databases must be registered in the recovery catalog before running an Oracle Backup job utilizing the **Record copies in RMAN recovery catalog** feature.
- In your Linux environment, if Oracle data or logs reside on LVM volumes, ensure the LVM version is 2.0.2.118 or later.
- For Oracle 12c databases, backups are created without placing the database in hot backup mode through Oracle Storage Snapshot Optimization. All associated snapshot functionality is supported. This feature requires the Advanced Compression feature of Oracle to be licensed. If this feature is not licensed in your environment, perform the following procedure to disable Snapshot Optimization and force the use of hot backup mode:

Create the file `/etc/guestapps.conf` on the Oracle server and add the following to it:

```
[DEFAULT]
skipSnapshotOptimization = true
```

If the file already exists, edit it and add the parameter under the existing [DEFAULT] section. This is a per-host setting. The parameter must be set in this file on each Oracle server where you want to force the use of hot backup mode.

- NOARCHIVELOG databases are not eligible for point-in-time recovery. NOARCHIVELOG databases can only be recovered to specific or latest versions. If upgrading from previous versions of IBM Storage Copy Data Management, the associated Oracle Inventory job must be re-run after upgrading to discover NOARCHIVELOG databases.
- When the option to create an additional log destination is selected, IBM Storage Copy Data Management automatically purges the logs under this new location after each successful backup. For IBM SVC, IBM Storage Copy Data Management purges logs after a FlashCopy operation but not after a Global Mirror operation. If both FlashCopy and Global Mirror are enabled for a database (whether in separate job definitions or the same), IBM Storage Copy Data Management purges the logs after the FlashCopy operation only. For databases that are protected only by a Global Mirror workflow, IBM Storage 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 Storage Copy Data Management does not purge logs from other log destinations so they must also be externally managed.
- If an Oracle Inventory job runs at the same time or short period after an Oracle Backup job runs, copy errors may occur due to temporary mounts that are created during the Backup job. As a best practice, schedule Oracle Inventory jobs so that they do not overlap with Oracle Backup jobs.

Considerations:

- Note that point-in-time recovery is not supported when one or more datafiles are added to the database in the period between the chosen point-in-time and the time that the preceding Backup job ran.

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

Important: You should not run Incremental FC and FC NoCopy backup jobs on the same source volume.

About this task

IBM Storage Copy Data Management provides application database copy management through application-consistent backup creation, cloning, and recovery. IBM Storage Copy Data Management copy management leverages the snapshot and replication features of the underlying storage platform to create, replicate, clone, and restore backups of Oracle databases. Archive log destinations as well as universal destination mount points are supported. Archived logs are automatically deleted upon reaching defined retention.

IBM Storage Copy Data Management auto-discovers databases and enables backups only of eligible databases. To be eligible for backup, application databases must reside on supported storage platforms.

The following options are available for Oracle Backup jobs:

RMAN Integration

Oracle Recovery Manager (RMAN), a command-line and Enterprise Manager-based tool, is the method preferred by Oracle database administrators for backup and recovery of Oracle databases, including maintaining an RMAN repository. The retention of RMAN cataloged data is managed by settings in Oracle. IBM Storage Copy Data Management automates cataloging of Oracle database backups in the RMAN recovery catalog, enabling database administrators to leverage RMAN for verification and advanced recovery.

Data Masking

Data masking is used to hide confidential data by replacing it with fictitious data. This feature is used when making data copies for DevTest or other use cases.

Log Backup

The log backup feature enables continuous backups of Archive logs to a specified destination.

Archive log retention is managed by settings in RMAN. IBM Storage Copy Data Management leverages archived logs to enable point-in-time recoveries of databases to facilitate RPOs.

Procedure

1. Click the **Jobs** tab. Expand the **Database** folder, then select **Oracle**.
2. Click **New**, then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more resources to back up. Expand Oracle home directories to view associated application databases.

Tip: You cannot select a database if it is not eligible for protection. Hover your cursor over the database name to view the reasons the database is ineligible, such as the database files, control files, or redo log files are stored on unsupported storage.

5. Select an SLA Policy that meets your backup data criteria.

Tip: For Sentinel scanning capabilities, the SLA policy must be Safeguarded Copy and you need to select a previously registered Security Scan server.

6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs** tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Tip: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum Concurrent Tasks

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

Skip IA Mount points and/or databases

Enable to skip Instant Disk Restore objects. By default, this option is enabled.

Record copies in RMAN local repository

Enable to create a local backup of the Recovery Manager (RMAN) catalog during the running of Oracle Backup job. RMAN catalogs can be used for backup, recovery, and maintenance of Oracle databases outside of IBM Storage Copy Data Management.

Record copies in RMAN recovery catalog

If Record copies in RMAN local repository is selected, select Record copies in RMAN recovery catalog to also create a remote RMAN catalog. Select an eligible Remote Catalog Database from the list of available sites. Select a Recovery Catalog Owner from the list of available Identities, or create a new Recovery Catalog Owner, then click OK.

Note that Oracle databases must be registered in the recovery catalog before running an Oracle Backup job utilizing the Record copies in RMAN recovery catalog feature.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#) on page 156.

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#)

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

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

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

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

- Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

- To edit Log Backup options before creating the job definition, click **Log Backup**. If **Create additional archive log destination** is selected, IBM Storage 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 Storage Copy Data Management automatically discovers the location where Oracle writes archived logs. If this location resides on storage from a supported vendor, IBM Storage 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 Storage 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.

Tip: NOARCHIVELOG databases are not eligible for log backup as they do not have archive logging enabled.

If multiple databases are selected for backup, then each of the servers hosting the databases must have their destination directories set individually. For example, if two databases from Server A and Server B are added to the same job definition, and a single destination directory named /logbackup is defined in the job definition, then you must create separate disks for both servers and mount them both to /logbackup on the individual servers.

If the **No archive logs / Use existing archive log destination(s)** option is selected, IBM Storage Copy Data Management does not automatically purge any archived logs. The retention of archived logs must be managed externally, for example using RMAN. In order to support point-in-time recovery, ensure that the retention period is at least large enough to retain all archived logs between successive runs of the Oracle Backup job.

If the **Create additional archive log destination** option is selected, IBM Storage Copy Data Management automatically manages the retention of only those archived logs that are under the new destination specified in the job definition. After a successful backup, logs older than that backup are automatically deleted from the IBM Storage Copy Data Management-managed destination. Even in this case, IBM Storage 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 Storage 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 Storage Copy Data Management job definition, the database parameter is not affected so if you want to stop using the log destination, you may need to manually disable it this parameter.

11. To edit Data Masking options before creating the job definition, click **Data Masking**. If enabled, IBM Storage Copy Data Management mounts snapshot copies of the protected database onto a user-specified staging server or source server. Select resources to be masked from the list of available databases, select a backup to mask, and an Oracle home where masking takes place. Set a trigger, then in the **Enter path to masking command on Oracle Server** field, enter the full path to an external script or tool to perform the data masking. For example, /home/oracle/tools/maskDatabase.sh.

Whether the masking takes place on the source server or a staging server, the masking process spins up the temporary database with a unique random name (for example, mask1234) that does not conflict with any other instance on that system. IBM Storage Copy Data Management then invokes the masking script with three arguments: the Oracle Home path, the new instance name, and the original instance name. For example:

```
/path/to/masking/script /u01/app/home1 mask1234 proddb
```

IBM Storage Copy Data Management spins up a clone of the database, then executes the user-specified command to perform masking. When the command completes successfully, IBM Storage 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 Storage Copy Data Management Restore jobs.

Tip: User-defined masking scripts that were in existence before IBM Storage Copy Data Management 2.2.7.4 must be updated to ensure they read the correct arguments and connect to the appropriate instance to perform masking.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.

What to do next

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

Related tasks

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

[“Creating an Oracle Restore job definition” on page 157](#)

The procedures describe how to create an Oracle Restore job definition.

Creating an Oracle application (on AIX) scanning backup job

This procedure describes how to create an Oracle-based application backup job definition for scanning for Sentinel.

Before you begin

- Create and run a Database Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - Database](#).
- Configure an SLA Policy. See [Configure SLA Policies](#).

- Register at least one AIX proxy host as file system in the application. When the user selects **Enable Proxy Server** all the AIX proxy host(s) that are registered in the IBM Storage Copy Data Management are displayed.
- Configure the NFS version 3 or 4 (preferably 4) on the AIX proxy and Security Scan server.
Note: If NFS mounting fails during the backup job execution, you must manually clean up the NFS mounts.
- A system property `protection.application.scan.fullNFSRestriction` has been introduced in `/opt/virgo/repository/ecx-usr/com.syncsort.dp.xsb.serviceprovider.properties`. By default the system property is set to **true** to create secure NFS for the Security Scan server, but if it is set to **false** IBM Storage Copy Data Management will create a global NFS share.
- Make sure that the proxy and scan host can ping each other's host addresses which are configured in IBM Storage Copy Data Management.
- Make sure that the proxy host is zoned and added to the IBM Storage array.

About this task

When creating a backup job of Oracle databases, IBM Storage Copy Data Management can utilize the IBM Storage Sentinel Anomaly Scan software to scan Oracle databases running on AIX for ransomware. An AIX proxy is utilized as part of this process.

Procedure

1. Click the **Jobs** tab. Expand the **Database** folder, then select **Oracle**.
2. Click **New**, then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more resources to back up.
5. Select an SLA Policy that meets your backup data criteria.

Tip: SLA policy must be Safeguarded Copy and previously registered Security Scan server needs to be selected.

6. Select **Enable Proxy Server** check box. All the AIX proxy host(s) that are registered in the IBM Storage Copy Data Management are displayed.
7. Choose the appropriate proxy host(s) for IBM Storage Sentinel Anomaly Scan software.
8. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs** tab. Repeat as necessary to add additional SLA Policies to the job definition. If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Tip: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

9. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.
10. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum Concurrent Tasks

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

Skip IA Mount points and/or databases

Enable to skip Instant Disk Restore objects. By default, this option is enabled.

Record copies in RMAN local repository

Enable to create a local backup of the Recovery Manager (RMAN) catalog during the running of Oracle Backup job. RMAN catalogs can be used for backup, recovery, and maintenance of Oracle databases outside of IBM Storage Copy Data Management.

Record copies in RMAN recovery catalog

If Record copies in RMAN local repository is selected, select Record copies in RMAN recovery catalog to also create a remote RMAN catalog. Select an eligible Remote Catalog Database from the list of available sites. Select a Recovery Catalog Owner from the list of available Identities, or create a new Recovery Catalog Owner, then click OK.

Note that Oracle databases must be registered in the recovery catalog before running an Oracle Backup job utilizing the Record copies in RMAN recovery catalog feature.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Scripts can also be uploaded and edited through the Scripts view on the **Configure** tab. For more information, see [Configure Scripts](#).

After the upload completes, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#)

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

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

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

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

11. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

12. To edit Log Backup options before creating the job definition, click **Log Backup**. If **Create additional archive log destination** is selected, IBM Storage 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 Storage Copy Data Management automatically discovers the location where Oracle writes archived logs. If this location resides on storage from a supported vendor, IBM Storage 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 Storage 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.

Tip: NOARCHIVELOG databases are not eligible for log backup as they do not have archive logging enabled.

If multiple databases are selected for backup, then each of the servers hosting the databases must have their destination directories set individually. For example, if two databases from Server A and Server B are added to the same job definition, and a single destination directory named /logbackup is defined in the job definition, then you must create separate disks for both servers and mount them both to /logbackup on the individual servers.

If the **No archive logs / Use existing archive log destination(s)** option is selected, IBM Storage Copy Data Management does not automatically purge any archived logs. The retention of archived logs must be managed externally, for example using RMAN. In order to support point-in-time recovery, ensure that the retention period is at least large enough to retain all archived logs between successive runs of the Oracle Backup job.

If the **Create additional archive log destination** option is selected, IBM Storage Copy Data Management automatically manages the retention of only those archived logs that are under the new destination specified in the job definition. After a successful backup, logs older than that backup are automatically deleted from the IBM Storage Copy Data Management-managed destination. Even in this case, IBM Storage 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 Storage 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 Storage Copy Data Management job definition, the database parameter is not affected so if you want to stop using the log destination, you may need to manually disable it this parameter.

13. To edit Data Masking options before creating the job definition, click **Data Masking**. If enabled, IBM Storage Copy Data Management mounts snapshot copies of the protected database onto a user-specified staging server or source server. Select resources to be masked from the list of available databases, select a backup to mask, and an Oracle home where masking takes place. Set a trigger, then in the **Enter path to masking command on Oracle Server** field, enter the full path to an external script or tool to perform the data masking. For example, /home/oracle/tools/maskDatabase.sh.

Whether the masking takes place on the source server or a staging server, the masking process spins up the temporary database with a unique random name (for example, mask1234) that does not conflict with any other instance on that system. IBM Storage Copy Data Management then invokes the masking script with three arguments: the Oracle Home path, the new instance name, and the original instance name. For example:

```
/path/to/masking/script /u01/app/home1 mask1234 proddb
```

IBM Storage Copy Data Management spins up a clone of the database, then executes the user-specified command to perform masking. When the command completes successfully, IBM Storage 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 Storage Copy Data Management Restore jobs.

Tip: User-defined masking scripts that were in existence before IBM Storage Copy Data Management 2.2.7.4 must be updated to ensure they read the correct arguments and connect to the appropriate instance to perform masking.

14. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.
15. You can customize the timeout value for the proxy server as per your need. By default the timeout for the proxy server is set to 600 seconds. Complete the following steps to change the timeout:
 - a. Go to the `/opt/virgo/repository/ecx-usr/com.syncsort.dp.xsb.serviceprovider.properties` path and change the value for the `protection.application.AIX.proxy.unmount.timeout.seconds` property.
 - b. Restart the virgo instance by issuing the following command:

```
systemctl restart virgo
```

For example:

Issue the following timeout for unmount command on AIX:

```
protection.application.AIX.proxy.unmount.timeout.seconds=<timeout_seconds>
```

where `<timeout_seconds>` represents the timeout value to enter.

Note: If you encounter the timeout issue, you may need to manually cleanup the stale disks from the proxy server and then proceed with the next job.

What to do next

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Hold, and Cancel a Job Session](#).
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#).
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Create an AIX-based Application Restore job definition. See [Create a Restore Job Definition - InterSystems Database](#).

Creating a SQL Backup job definition

The procedures describe how to create a SQL Backup job definition.

Before you begin

- Create and run a Database Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - Database](#).
- Configure an SLA Policy. See [Configure SLA Policies](#).
- Review SQL requirements. See [Microsoft SQL requirements](#).

Microsoft SQL server considerations:

- UUID must be enabled to perform Microsoft SQL-based backup functions. To enable, power off the guest machine through the vSphere client, then select the guest and click **Edit Settings**. Select **Options**, then **General** under the Advanced section. Select **Configuration Parameters...**, then find the disk.

EnableUUID parameter. If set to FALSE, change the value to TRUE. If the parameter is not available, add it by clicking **Add Row**, set the value to TRUE, then power on the guest.

- You must manually create a directory to store VSS provider logs when running IBM Storage Copy Data Management 2.2.6 and earlier. Create the following directory structure on the SQL server:
c:\temp\CDM\logs
- During a Backup job, an error may display in the logs stating that the size of the manifest file is a specific size, and the snapshot is considered not application consistent. If this error displays, it may be due to the number of disks associated with a particular SCSI controller. As a best practice, configure your disks so that all controllers have at least half of their slots empty.
- IBM Storage Copy Data Management supports one Microsoft SQL database per SQL Backup job, therefore you must avoid protecting a SQL database through multiple Backup jobs.
- Note that IBM Storage Copy Data Management does not support log backup of Simple recovery models.
- An AlwaysOn of a replica of a SQL cluster instance is not supported. Replicas are limited standalone SQL servers and instances.
- **Considerations:**
 - Note that point-in-time recovery is not supported when one or more datafiles are added to the database in the period between the chosen point-in-time and the time that the preceding Backup job ran.
 - For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).
 - One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#).

Important: You should not run Incremental FC and FC NoCopy backup jobs on the same source volume.

About this task

IBM Storage Copy Data Management provides application database copy management through application-consistent backup creation, cloning, and recovery. IBM Storage Copy Data Management leverages the snapshot and replication features of the underlying storage platform to create, replicate, clone, and restore copies of Microsoft SQL Servers. Archive log destinations as well as universal destination mount points are supported. Archived logs are automatically deleted upon reaching defined retention.

IBM Storage Copy Data Management auto-discovers databases and enables backups only of eligible databases. To be eligible for backup, application databases must reside on supported storage platforms.

The following options are available for SQL Backup jobs:

Log Backup

The log backup feature enables continuous backup of Archive logs to a specified destination.

IBM Storage Copy Data Management leverages archived logs to enable point-in-time recoveries of databases to facilitate RPOs.

Procedure

1. Click the **Jobs** tab. Expand the **Database** folder, then select **SQL**.
2. Click **New**, then select **Backup**. The job editor opens.
3. Select a **Standalone** or **Failover Cluster** or **Always On Availability Group** workflow template.
4. Enter a name for your job definition and a meaningful description.
5. From the list of available sites select one or more resources to back up. Expand servers to view associated application databases.

Tip: You cannot select a database if it is not eligible for protection. Hover your cursor over the database name to view the reasons the database is ineligible, such as the database files, control files, or redo log files are stored on unsupported storage.

6. Select an SLA Policy that meets your backup data criteria.

7. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs** tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Tip: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

8. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum Concurrent Tasks

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

Skip IA Mount points and/or databases

Enable to skip Instant Disk Restore objects. By default, this option is enabled.

Maximum Concurrent Snapshots on ESX

Set the maximum number of concurrent snapshots on the vCenter.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

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

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the

job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

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

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

11. To edit Log Backup options before creating the job definition, click **Log Backup**. If **Backup Logs** is selected, IBM Storage Copy Data Management backs up database logs then protects the underlying disks. Select resources in the **Select resource(s) to add archive log destination** field. Database logs are backed up to the directory entered in the **Use Universal Destination Mount Point** field, or in the Mount Point field after resources are selected. The destination must already exist and must reside on storage from a supported vendor.

Tip:

Always On job definitions must specify the log destination as a path in the following format: \
\server\share\optional_subfolder. The server can be either an IP address or hostname that is resolvable from the IBM Storage Copy Data Management appliance.

If multiple databases are selected for backup, then each of the servers hosting the databases must have their Destination Mount Points set individually. For example, if two databases, one from Server A and one from Server B, are added to the same job definition, and a single mount point named /logbackup is defined in the job definition, then you must create separate disks for each server and mount them both to /logbackup on the individual servers.

IBM Storage Copy Data Management automatically truncates post log backups of databases that it backs up. If database logs are not backed up with IBM Storage Copy Data Management, logs are not truncated by IBM Storage Copy Data Management and must be managed separately.

To disable a log backup schedule on the SQL server, edit the associated SQL Backup job definition and deselect the checkbox next to the database on which you wish to disable the log backup schedule in the **Select resource(s) for log backup destination** field, then save and re-run the job. Note that the job definition must be saved and re-run for the disablement to take effect.

When SQL backup job completes with log backups enabled, all transaction logs up to the point of the job completing are purged from the SQL server. Note that log purging will only occur if the SQL Backup job completes successfully. If log backups are disabled during a re-run of the job, log purging will not occur.

If a source database is overwritten, all old transaction logs up to that point are placed in a “condense” directory once the restoration of the original database completes. When the next run of the SQL Backup job completes, the contents of the condense folder is removed.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.

What to do next

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

Related tasks

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

[“Creating a Microsoft SQL Restore job definition” on page 162](#)

The procedures describe how to create a Microsoft SQL Restore job definition.

Creating a File System Backup job

The procedures describe how to create a File System Backup job definition.

Before you begin

- Review File System requirements. See [File system requirements](#).
- Create and run a File System Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - File System](#).
- Configure an SLA Policy. See [Configure SLA Policies](#).

Considerations:

- An AlwaysOn of a replica of a SQL cluster instance is not supported. Replicas are limited standalone SQL servers and instances.
- Note that point-in-time recovery is not supported when one or more datafiles are added to the database in the period between the chosen point-in-time and the time that the preceding Backup job ran.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#).

Important: You should not run Incremental FC and FC NoCopy backup jobs on the same source volume.

About this task

IBM Storage Copy Data Management provides application database copy management of physical Windows, Linux, and AIX file systems through a File System Backup job.

Procedure

1. Click the **Jobs** tab. Expand the **Database** folder, then select **File System**.
2. Click **New**, then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more resources to back up.

Tip: You cannot select a database if it is not eligible for protection. Hover your cursor over the database name to view the reasons the database is ineligible, such as the database files, control files, or redo log files are stored on unsupported storage.

5. Select an SLA Policy that meets your backup data criteria.
6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs** tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Tip: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum Concurrent Tasks

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

Skip IA Mount points and/or databases

Enable to skip Instant Disk Restore objects. By default, this option is enabled.

Maximum Concurrent Snapshots on ESX

Set the maximum number of concurrent snapshots on the vCenter.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

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

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

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

Note: If adding a script to a Windows-based File System job definition, the user running the script must have the "Log on as a service" right enabled, which is required for running prescripts and postscripts. For more information about the "Log on as a service" right, see [Add the Log on as a service Right to an Account](#).

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

10. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.

What to do next

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#).
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#).
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Create an OS Volume Restore job definition. See [Create a Restore Job Definition - File System](#).

Related tasks

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

[“Creating a File System Restore job definition” on page 166](#)

The procedures describe how to create a File System Restore job definition.

Creating an IBM Storage Virtualize Backup job definition

The procedures describe how to create an IBM Storage Virtualize Backup job definition.

Before you begin

- Create and run an IBM Storage Virtualize Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - IBM Storage Virtualize](#).
- Configure an SLA Policy. See [Configure SLA Policies](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

- IBM providers utilize port 22 for communication with IBM Storage Copy Data Management.
- Note that snapshot postscript functionality applies only to FlashCopy subpolicies.
- In IBM storage environments, port grouping and IP partnerships are required to enable remote copy connections. See IBM's SAN Volume Controller and Storwize Family Native IP Replication Guide.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#).

Important: You should not run Incremental FC and FC NoCopy backup jobs on the same source volume.

About this task

Back up IBM data with FlashCopies, Global Mirrors, and VM Copies using an IBM Storage Virtualize Backup job. The RPO and copy data parameters are defined in an SLA Policy, which is then applied to the Backup job definition, along with a specified activation time to meet your copy data criteria. Supported sources include IBM Storage Virtualize storage systems.

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **IBM Storage Virtualize**.
2. Click **New**, then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more providers to back up. To exclude Flash Copies from the list of sources, select **Exclude Flash Copies**.
5. Select an SLA Policy that meets your backup data criteria.
6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs** tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Tip: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Skip the Flash Copy Target Volumes

Select this option to ensure FlashCopy target volumes are excluded from jobs associated with the SLA Policy.

Maximum Concurrent Tasks

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

Create Consistency Group

If multiple volumes are selected in the Source tab (for example, volumes that contain data tied to an application) enable this option to add the volumes to a FlashCopy or Global Mirror Consistency Group to perform Copy Data functions on the entire group. If the associated SLA Policy contains both FlashCopy and Global Mirror subpolicies, a separate Consistency Group will be created for each copy type. Note that if more than one IBM provider is selected in the job definition, a Consistency Group will be created for each provider. Consistency Groups are named based on the prefix provided during job creation plus the job name.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

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

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the Configure tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

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

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

10. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.

What to do next

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#).
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#).
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Use the Inventory Browse feature to review the recovery point. See [Browse Inventory](#).
- Create an IBM Storage Virtualize Restore job definition. See [Create a Restore Job Definition - IBM Storage Virtualize](#).

Related concepts

[“Creating an IBM Storage Virtualize Restore job definition” on page 169](#)

Related tasks

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Creating a Safeguarded Copy Backup

This procedure describes how to create a Safeguarded Copy (SGC) backup.

Before you begin

- Create and run a Safeguarded Copy that you want to backup. For more information, see [Creating a Safeguarded Copy SLA policy](#).
- Configure an SLA Policy. See [Configure SLA Policies](#).

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **IBM Storage Virtualize**.
2. Click **New**, then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available volumes select the volume to backup. To exclude Flash Copies from the list of sources, select **Exclude Flash Copies**.
5. Select an SLA Policy that meets your backup data criteria.
6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs** tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Tip: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Skip the Flash Copy Target Volumes

Select this option to ensure FlashCopy target volumes are excluded from jobs associated with the SLA Policy.

Maximum Concurrent Tasks

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

Create Consistency Group

If multiple volumes are selected in the Source tab (for example, volumes that contain data tied to an application) enable this option to add the volumes to a FlashCopy or Global Mirror Consistency Group to perform Copy Data functions on the entire group. If the associated SLA Policy contains both FlashCopy and Global Mirror subpolicies, a separate Consistency Group will be created for each copy type. If more than one IBM provider is selected in the job definition, a Consistency Group is created for each provider. Consistency Groups are named based on the prefix that is provided during job creation plus the job name.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

After completion, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

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

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

After completion, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials that are required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots that are created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

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

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

10. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.

What to do next

Note:

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#).
- Track the progress of the job session on the **Jobs** tab. See [Monitor a Job Session](#).

- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Use the Inventory Browse feature to review the recovery point. See [Browse Inventory](#).
- Create a Safeguarded Copy Restore. See [Creating a Safeguarded Copy Restore](#).

Creating an IBM Storage Virtualize for Snapshot Backup job definition

The procedures describe how to create an IBM Storage Virtualize for Snapshot Backup job definition.

Before you begin

- Create and run an IBM Storage Virtualize for Snapshot Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - IBM Storage Virtualize](#).
- Configure an SLA Policy. See [Configure SLA Policies](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

- IBM providers utilize port 22 for communication with IBM Storage Copy Data Management.
- Note that snapshot postscript functionality applies only to Snapshot subpolicies.
- In IBM storage environments, port grouping and IP partnerships are required to enable remote copy connections. See IBM's SAN Volume Controller and Storwize Family Native IP Replication Guide.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#).

Important: You should not run Incremental FC and FC NoCopy backup jobs on the same source volume.

About this task

Back up IBM data with Snapshots, Global Mirrors, and VM Copies using an IBM Storage Virtualize for Snapshot Backup job. The RPO and copy data parameters are defined in an SLA Policy, which is then applied to the Backup job definition, along with a specified activation time to meet your copy data criteria. Supported sources include IBM Storage Virtualize for Snapshot storage systems.

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **IBM Storage Virtualize for Snapshot**.
2. Click **New**, then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more providers to back up. To exclude Flash Copies from the list of sources, select **Exclude Flash Copies**.
5. Select an SLA Policy that meets your backup data criteria.
6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs** tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Tip: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Skip the Flash Copy Target Volumes

Select this option to ensure FlashCopy target volumes are excluded from jobs associated with the SLA Policy.

Maximum Concurrent Tasks

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

Create Consistency Group

If multiple volumes are selected in the Source tab (for example, volumes that contain data tied to an application) enable this option to add the volumes to a FlashCopy or Global Mirror Consistency Group to perform Copy Data functions on the entire group. If the associated SLA Policy contains both FlashCopy and Global Mirror subpolicies, a separate Consistency Group will be created for each copy type. Note that if more than one IBM provider is selected in the job definition, a Consistency Group will be created for each provider. Consistency Groups are named based on the prefix provided during job creation plus the job name.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

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

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

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

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

10. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.

What to do next

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#).
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#).
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Use the Inventory Browse feature to review the recovery point. See [Browse Inventory](#).
- Create an IBM Storage Virtualize Restore job definition. See [Create a Restore Job Definition - IBM Storage Virtualize](#).

Creating a Safeguarded Copy Backup

This procedure describes how to create a Safeguarded Copy (SGC) backup.

Before you begin

- Create and run a Safeguarded Copy that you want to backup. For more information, see [Creating a Safeguarded Copy SLA policy](#).
- Configure an SLA Policy. See [Configure SLA Policies](#).

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **IBM Storage Virtualize for Snapshot**.
2. Click **New**, then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available volumes select the volume to backup. To exclude Flash Copies from the list of sources, select **Exclude Flash Copies**.
5. Select an SLA Policy that meets your backup data criteria.
6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs** tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Tip: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Skip the Flash Copy Target Volumes

Select this option to ensure FlashCopy target volumes are excluded from jobs associated with the SLA Policy.

Maximum Concurrent Tasks

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

Create Consistency Group

If multiple volumes are selected in the Source tab (for example, volumes that contain data tied to an application) enable this option to add the volumes to a Snapshot to perform Copy Data functions on the entire group.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

After completion, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

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

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

After completion, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials that are required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots that are created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

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

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

10. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.

What to do next

Note:

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

Creating a NetApp ONTAP Backup job definition

The procedures describe how to create a NetApp ONTAP Backup job definition.

Before you begin

- Create and run a NetApp ONTAP Storage Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - NetApp ONTAP Storage](#).
- Configure an SLA Policy. See [Configure SLA Policies](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

- Note that NetApp ONTAP Backup jobs can only vault or mirror snapshots created through IBM Storage Copy Data Management jobs.
- Note that cloned volumes will not be replicated through a Backup job.
- Note that snapshot postscript functionality applies only to NetApp ONTAP storage snapshot subpolicies.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#).

Important: You should not run Incremental FC and FC NoCopy backup jobs on the same source volume.

About this task

Back up NetApp ONTAP data using a NetApp ONTAP Backup job. A NetApp ONTAP Backup job consists of snapshot, mirror, and vault sub-policies defined in an SLA Policy, each with their own set of options to give you more control of your NetApp ONTAP protection needs. After an initial primary snapshot is added to the workflow, additional vaults and mirrors ensure your data is replicated to multiple locations.

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **NetApp ONTAP**.
2. Click **New**, then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more providers to back up.

5. Select an SLA Policy that meets your backup data criteria.
6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs** tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Tip: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum Concurrent Tasks

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

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

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

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

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

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

10. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.

What to do next

- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#).
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#).
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.
- Use the Inventory Browse feature to review the recovery point. See [Browse Inventory](#).
- Create a NetApp ONTAP Restore job definition. See [Create a Restore Job Definition - NetApp ONTAP](#).

Related concepts

[“Creating an NetAPP ONTAP Restore job definition” on page 179](#)

Related tasks

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Creating a Pure Storage FlashArray Backup job definition

The procedures describe how to create a Pure Storage FlashArray Backup job definition.

Before you begin

- Create and run a Pure Storage FlashArray Inventory job that includes the providers you wish to back up. See [Create an Inventory Job Definition - Pure Storage FlashArray](#).
- Configure an SLA Policy. See [Configure SLA Policies](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

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

Important: You should not run Incremental FC and FC NoCopy backup jobs on the same source volume.

About this task

Back up Pure Storage data with snapshots and replications using a Pure Storage FlashArray Backup job. The RPO and copy data parameters are defined in an SLA Policy, which is then applied to the Backup job definition along with a specified activation time to meet your copy data criteria.

Pure Storage FlashArray can be set to utilize CloudSnap functionality so that snapshots are offloaded from the local storage array to cloud storage. If the Pure Storage SLA was created with the Add Snapshot Offload option, a copy of the snapshot on the local storage array will be offloaded to a S3 cloud storage target or NFS share and adhere to the number of snapshots or the number of days as provided when creating the sub-policy.

Tip: The IBM Storage Copy Data Management user interface will indicate that a Pure Storage FlashArray CloudSnap offload job has completed even though the transfer is still occurring in the background which is dependent on network speeds. Consider setting an age as the retention for offload copies when using the Pure Storage FlashArray CloudSnap functionality. Doing so will ensure that a sufficient amount of time has passed for data to be transferred to the S3 storage target or NFS share before it has to be condensed out from a backup. This is particularly important if several offload jobs are run in quick succession.

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **Pure Storage FlashArray**.

Tip: Pure Storage FlashArray requires that the consistency group option be selected. Click **Advanced** to enable this option.

2. Click **New**, then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available sites select one or more providers to back up.
5. Select an SLA Policy that meets your backup data criteria.
6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs** tab. Repeat as necessary to add additional SLA Policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA Policies to run concurrently.

Tip: Only SLA Policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum Concurrent Tasks

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

Create Consistency Group

If multiple volumes are selected in the Source tab (for example, volumes that contain data tied to an application) enable this option to add the volumes to a Consistency Group to perform backup functions on the entire group. If the associated SLA Policy contains different subpolicy types, a separate Consistency Group will be created for each backup type. Consistency Groups are named based on the prefix provided during job creation plus the job name.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

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

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the Configure tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

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

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

10. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.

What to do next

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

- Create a Pure Storage FlashArray Restore job definition. See [Create a Restore Job Definition - Pure Storage FlashArray](#).

Related concepts

[“Creating a Pure Storage FlashArray Restore job definition” on page 187](#)

Related tasks

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Creating a VMware Backup job definition

The procedures describe how to create a VMware Backup job definition.

Before you begin

- Create and run a VMware Inventory job that includes the providers you want to back up. See [Create an Inventory Job Definition - VMware](#).
- Configure an SLA Policy. See [Configure SLA Policies](#).
- Ensure the latest version of VMware Tools is installed in your environment. IBM Storage Copy Data Management was tested against VMware Tools 9.10.0.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

- VMware Backup and Restore jobs support only vCenters or ESX hosts running vSphere 6.0 through 7.0.
- When running VADP-based VM Replication workflows, target volumes and datastores can be automatically expanded in response to space usage requirements if supported by the underlying storage. Automatic growing prevents a volume from running out of space or forcing you to delete files manually. For a list of supported storage systems, see [System requirements](#).
- VMware DRS cluster datastores are supported in VMware Backup and Restore jobs.
- In NetApp ONTAP environments running Clustered Data ONTAP, cluster peering must be enabled. Peer relationships enable communication between SVMs. See NetApp ONTAP's Cluster and Vserver Peering Express® Guide.
- In addition to NFS, IBM Storage 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 that are associated with a VMware provider must be added to IBM Storage Copy Data Management, which include NetApp ONTAP storage controllers and clusters. See [Register a Provider](#).
- VMware Backup jobs do not support virtual machine SCSI controllers where the SCSI Bus Sharing value is set to virtual or physical.
- Instant Disk Restore recoveries that use the VM Replication method are not supported at the datastore level. Instant Disk Restore datastore level recoveries are supported through the primary storage snapshot method.
- Snapshot protection is not supported at an ESX server level.
- Cloned volumes will not be replicated through a backup job.

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

Important: Do not run Incremental FC and FC NoCopy backup jobs on the same source volume.

Considerations for VMware virtual volumes:

- SLA policies that include virtual machines that are stored on virtual volume (VVOL) datastores through VM Replication sub-policies are supported. Replication is supported on the VM Replication target.
- Storage snapshots of virtual machines that reside on a VVOL are currently not supported. If a storage snapshot operation is selected for a virtual machine that resides on a VVOL, the virtual machine is skipped.

Autogrow for VM replication (VADP) requirements:

- When running VADP-based VM Replication-based SLA policies, target volumes and datastores can be automatically expanded in response to space usage requirements if supported by the underlying storage. The following storage systems are supported:
 - NetApp ONTAP Clustered Data ONTAP 8.2 and later (including 7-mode) are supported for all SLA Policy types. Versions earlier than 8.2 are supported for all SLA Policy types except Replication.
 - IBM storage systems running IBM Storage Virtualize software are supported for all SLA Policy types except Replication.
 - Pure Storage systems are supported for all SLA Policy types.

About this task

Back up VMware data including virtual machines, datastores, folders, vApps, and data centers with snapshots by using a VMware Backup job. A VMware Backup job consists of snapshot, mirror, and vault sub-policies that are defined in an SLA Policy, each with their own set of options to give you more control of your VMware protection needs. After an initial primary snapshot is added to the workflow, additional vaults and mirrors ensure that your data is replicated to multiple locations.

VMware Backups jobs support IBM, NetApp ONTAP, and Pure Storage FlashArray SLA policies.

Procedure

1. Click the **Jobs** tab. Expand the **Hypervisor** folder, then select **VMware**.
2. Click **New**, then select **Backup**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the drop-down menu, select **VMs and Templates or Storage**. From the list of available sites, select one or more resources to back up, including virtual machines, VM templates, datastores, folders, vApps, and data centers.
5. Select an SLA Policy that meets your backup data criteria.

Tip: For Sentinel scanning capabilities, the SLA policy must be Safeguarded Copy and you need to select a previously registered Security Scan server.

6. Click the job definition's associated **Schedule Time** field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on demand through the **Jobs** tab. Repeat as necessary to add additional SLA policies to the job definition.

If configuring more than one SLA Policy in a job definition, select the **Same as workflow** option to trigger multiple SLA policies to run concurrently.

Tip: Only SLA policies with the same RPO frequencies can be linked through the **Same as workflow** option. Define an RPO frequency when creating an SLA Policy.

7. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.
8. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Maximum Concurrent Tasks

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

Create VM snapshots for all VMs

Enable to configure virtual machine snapshot options. Available options include creating virtual machine snapshots for all virtual machines, making all virtual machines included in the job application or file system consistent, or making specific virtual machines included in the job application or file system consistent.

Application consistent backup data captures data in memory and transactions in process. All VSS-compliant applications such as Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, Microsoft SQL, and system state are quiesced. VMDKs and virtual machines can be instantly mounted to recover data related to quiesced applications.

Truncate application logs

To truncate application logs for SQL during the Backup job, enable the **Truncate application logs** option. Note that credentials must be established for the associated virtual machine and SQL instance through the **Sites & Providers** pane on the **Configure** tab. Select a VMware provider, click the **VMs** tab, then click the associated virtual machine. Click the **Credentials** tab and add credentials for the virtual machine. Note that System credentials are always required. If the credentials are the same for the SQL instance, select the **Use System Credentials for app** option. If the credentials differ, you must provide credentials for all SQL instances, including the default SQL server. Ensure that the **Type** field in the **New Credential** dialog window is set to **SQL**.

IBM Storage Copy Data Management generates logs pertaining to the application log truncation function and copies them to the following location on the IBM Storage Copy Data Management appliance: `/data/log/ecxdeployer/<vm name>/logs`.

VM Snapshot Scripts

VM snapshot prescripts and postscripts are scripts that can be run on the virtual machine before or after a VMware virtual machine snapshot is taken. The snapshot prescript runs before a VMware virtual machine snapshot is captured, while the snapshot postscript runs after the snapshot completes. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines. See [Configure Scripts](#).

Select a virtual machine, then click the Scripts field in the Pre-Script or Post-Script section to select or upload a script. Once complete, the script displays in the Selected Script(s) section. Click the Parameters field to add a parameter to the script, then click Add. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking Add.

Click the Identity field to add or create the credentials required to run the script. See [Identities Overview](#).

Repeat the procedure for each virtual machine associated with the job.

Tip: When the VM Snapshot script is run, the virtual machine name will be passed as the first argument to the script. Any additional arguments specified in the job will follow as second, third, and so forth. If a non-zero exit code is returned by the script, the associated snapshot task fails.

Skip read only datastores

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

Skip IA Mount points and/or databases

Enable to skip Instant Disk Restore objects. By default, this option is enabled.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

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

Enable Job-level Snapshot Scripts

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot task runs. The snapshot prescript runs before all associated snapshots are run, while the snapshot postscript runs after all associated snapshots complete. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the Configure tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed. For parameter examples, see [Using State and Status Arguments in Postscripts](#).

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

SNAPSHOTS is an optional parameter for snapshot postscripts that displays a comma separated value string containing all of the storage-based snapshots created by the job. The format of each value is as follows: <registered provider name>:<volume name>:<snapshot name>.

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

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

10. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.

What to do next

- If in a Linux environment, consider creating VADP proxies to enable load sharing. See [Create VMware Backup Job Proxies](#).
- If you do not want to wait until the next scheduled job run, run the job session on demand. See [Start, Pause, and Hold a Job Session](#).
- Track the progress of the job session on the Jobs tab. See [Monitor a Job Session](#).
- If notification options are enabled, an email message with information about the status of each task is sent when the job completes.

- Use the Inventory Browse feature to review the recovery point. See [Browse Inventory](#).
- Create a VMware Restore job definition. See [Create a Restore Job Definition - VMware](#).

Related concepts

[“Configure scripts” on page 85](#)

[“Creating VMware Backup job proxies” on page 149](#)

[“Creating a VMware Restore job definition” on page 191](#)

Related tasks

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Creating VMware Backup job proxies

Before you begin

- Determine how many proxies to create; the more proxies, the faster the jobs run. Each proxy is used merely to process data, and the results are returned through the host IBM Storage Copy Data Management server. The only impact seen by the user is the improvement in performance when running the job.

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

Note that the first time a given job is run, the proxies do not take effect because VM clone technology is used. But the second and subsequent times the job is run, change block tracking technology is used and the proxies are employed.

If proxies exist, the entire processing load is shifted off the IBM Storage Copy Data Management host machine and onto the proxies, else the entire load stays on the IBM Storage Copy Data Management host. Within a Backup job, the processing load for any single VM is shifted to a single proxy machine; multiple VMs are shifted to multiple proxies if available.

If a proxy server goes down or is otherwise disabled before the start of the job, the other proxies (or if there are no other proxies, the IBM Storage 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.

System Requirements:

This feature has been tested only for Ubuntu, SUSE Linux Enterprise Server, and Red Hat® environments. It is supported only in x64 configurations with a minimum kernel of 2.6.32.

A minimum of 8 GB of RAM is required (16 GB recommended), along with 60 GB of disk space.

Each proxy must have a fully qualified domain name.

Installer Notes:

The IBM Storage Copy Data Management 2.2.6+ of the VADP Proxy installer includes Virtual Disk Development Kit (VDDK) 5.5.5. This version of the VADP proxy installer provides the following functionality:

- External VADP Proxy support with vSphere 6.5.
- External VADP Proxy support for Hot Add operations, which provides higher performance for VADP Backups.

Creating 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 Storage Copy Data Management machine.
2. Copy the VADP Proxy installation program to the local proxy machine. The VADP Proxy file can be downloaded from the MySupport site.
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.

Tip: 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 Storage Copy Data Management host.
 - a. When prompted for the installation directory, select **/opt/CDM**.
 - b. When prompted for the IBM Storage Copy Data Management Discovery Server IP, enter the IP Address of the IBM Storage Copy Data Management host.
 - c. When prompted for the IBM Storage 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 `ecxvdp` is started on the proxy machine. A log file `ecxvdp.log` is generated in `/opt/CDM/logs` directory.

7. Repeat the previous steps for each proxy you want to create.

Confirming a proxy server connections:

1. Start the IBM Storage Copy Data Management management console.
2. In the IBM Storage Copy Data Management management console, click the arrow next to the **Support** icon, then choose **View Edge Services Status**. The IBM Storage Copy Data Management host machine and each operational proxy server are displayed.

Next steps

- Run the VMware Backup job. The use of the proxies are indicated in the job log by a log message similar to the following:

```
Run remote vmdkbackup of MicroService: http://<proxy node name>, IP:<proxy IP address>
```

- Uninstall the proxies when you cease running the VMware Backup jobs. To uninstall a proxy, on your host machine, run the following command from the `uninstall` subdirectory of the installation directory (default installation directory is `/opt/CDM`):

```
./uninstall_vmdkbackup
```

The installation directory is removed.

Related concepts

[“Start, pause, and hold a job session” on page 91](#)

From the Jobs tab you can run a job session on demand, pause or cancel a running job, and hold all future scheduled instances of a job from running until you are ready for the job to proceed.

Related tasks

[“Creating a VMware Backup job definition” on page 145](#)

The procedures describe how to create a VMware Backup job definition.

Restore jobs

The topics in the following section cover Restore job definitions as well as postscript argument details.

Creating an InterSystems Database Restore job definition

The procedures describe how to create an InterSystems Database Restore job definition. You can create InterSystems Caché and InterSystems IRIS application servers in IBM Storage Copy Data Management as an InterSystems Database.

Before you begin

- Create and run an InterSystems Database Backup job. See [Create a Backup Job Definition - InterSystems Database](#).
- Review InterSystems Database requirements. See [InterSystems database requirements](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).
- You must add credentials to the destination virtual machine when recovering with the subnet option. See [Add Credentials to a Virtual Machine](#).

Considerations:

- Note that the following users and groups must be created on the target host: instance owner, effective user for InterSystems Database superserver and its jobs, effective group for InterSystems Database processes, and a group that has permissions to start and stop InterSystems Database instances. The user and group IDs should match those on the source host. The instance will be brought up using the same mount points as those found on the source machine, so ensure these mounts are not in use on the target.
- Note that it is possible to scan in an InterSystems Database backup failover member instance or an async member instance and run snapshots against the mirror copy instead of the primary failover member.
- When creating an InterSystems Database restore job definition, select only one instance to restore. If more than once instance is selected, the InterSystems Database agent only restores the last instance it receives in the command request.
- When restoring to a target with running InterSystems Database instances, the instances display as valid targets. Note that IBM Storage Copy Data Management will not interact with these instances, but instead bring up a new instance using mapped mount points. When restoring to a target with no prior InterSystems Database instances, IBM Storage Copy Data Management creates a placeholder that acts as a restore target named `cache_general`. Note that `cache_general` should only be used as a restore target and should not be selected for backup.
- Single InterSystems Database databases can be restored through an Instant Disk Restore job, which mounts physical volumes on the target machine. Granular recovery can then be performed through InterSystems Database commands.

About this task

IBM Storage Copy Data Management leverages Copy Data Management technology for recovering application databases through Database Restore jobs. Your InterSystems Database clones can be utilized and consumed instantly through IBM Storage Copy Data Management Instant Disk Restore jobs. IBM

Storage Copy Data Management catalogs and tracks all cloned instances. Instant Access leverages iSCSI or fibre channel protocols to provide immediate mount of LUNs without transferring data. Snapshotted databases are cataloged and instantly recoverable with no physical transfer of data.

Tip: Create a schedule before creating a job definition so that you can easily add the schedule to the job definition.

Procedure

1. Click the **Jobs** tab. Expand the **Database** folder, then select **InterSystems Database**.
2. Click **New**, then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select a template. Available options include **Instant Database Restore** and **Instant Disk Restore**.
5. Click **Source**. From the drop-down menu select **Application Browse** to select a source site and an application server to view available database recovery points. Select resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **Application Search** from the drop-down menu to search for application servers with available recovery points. Add copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy**. Sites containing copies of the selected data display. Select a site. By default, the latest copy will be used. If you would like to choose a specific version or use latest successful scan, select a site then click the **Select Specific Version or Use Latest Successful Scan**. Click the **Use Latest** or **Use Latest Successful Scan** to view specific copies and their associated job and completion time. If a selected recovery of a snapshot fails, another selection for the recovery should be done manually in the restore job, which will be a copy from the same site being used.
7. Click **Destination**. Select a source site and an associated destination. If creating an **Instant Disk Restore** job definition, review the destination's database name mapping settings. Optionally, click the **New database name** field to create an alternate database name.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Application Options

Rename Mount Points

For more information about the **Rename Mount Points** options, see [Restore Jobs - Rename Mount Points and Initialization Parameter Options](#).

If creating an Instant Database Restore job definition, this option is set to Do Not Rename by default. IBM Storage Copy Data Management will mount the mount points with the same path/name as the source.

Change the GUID of the Instance

By default the restored instance will have the same GUID as the source instance. If creating an Instant Database Restore job definition, selecting this option will generate a new GUID to be assigned to the restored instance.

Policy Options

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If unselected, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the database recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Allow to overwrite vDisk

In cases where the Make Permanent option is enabled, and the destination VM has conflicting VMDK files, enable the **Allow to overwrite vDisk** option to delete the existing VMDK and overwrite it with the selected source.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#). **State** and **status** arguments are not supported for Backup jobs.

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

Storage Options

Make Permanent

Set the default permanent restoration action of the job. All database recovery operations can leverage Instant or Test modes and then either be deleted or promoted to permanent mode. This behavior is controlled through the **Make Permanent** option.

Important: When you run a restore with **Make Permanent**, the data is vMoted to the datastore where the virtual machine's (VM) VMX file resides. The datastore to which the data and logs are moved may not be ideal or even supported by IBM Storage Copy Data Management for future backups. In most cases, the data and logs will be moved to the same datastore as the VM operating system disk and this can result in subsequent backup failures. You should inspect the VM configuration after restore with **Make Permanent** completes and manually reconfigure the VM to move the data and logs disks to datastores that are supported for subsequent backups and not the datastore containing the VM operating system disk. Then run VM and application inventory jobs explicitly to capture the updated configuration for the application servers. Finally, you can run another backup job of the resource so that a snapshot is available for future restore jobs.

Enabled - Always make permanent through full copy FlashCopy

Disabled - Never make permanent

User Selection - Allows the user to select Make Permanent or Cleanup when the job session is pending

Protocol Priority

If more than one storage network protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

What to do next

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

Related tasks

[“Creating an InterSystems Database Backup job definition” on page 110](#)

The procedures describe how to create an InterSystems Database Backup job definition. You can create InterSystems Caché and InterSystems IRIS application servers in IBM Storage Copy Data Management as an InterSystems Database.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Creating a SAP HANA Restore job definition

This procedure describes how to create a SAP HANA Restore job definition.

Before you begin

- Create and run a SAP HANA Backup job. See [Create a Backup Job Definition - SAP HANA](#).
- Review SAP HANA requirements. See [SAP HANA requirements](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).
- You must add credentials to the destination virtual machine when recovering with the subnet option. See [Add Credentials to a Virtual Machine](#).
- When performing a SAP HANA Disk Restore, point-in-time (PIT) is an option. For restore jobs of SAP HANA log backups, PIT cannot be used. Instead, restore at the snapshot level and use SAP HANA Studio to do a PIT restore.

About this task

IBM Storage Copy Data Management leverages Copy Data Management technology for recovering application databases through Database Restore jobs. Your SAP HANA clones can be utilized and consumed instantly through IBM Storage Copy Data Management Instant Disk Restore jobs. IBM Storage Copy Data Management catalogs and tracks all cloned instances. Instant Disk Restore leverages ISCSI or fibre channel protocols to provide immediate mount of LUNs without transferring data. Snapshotted databases are cataloged and instantly recoverable with no physical transfer of data.

Tip: Create a schedule before creating a job definition so that you can easily add the schedule to the job definition.

Procedure

1. Click the **Jobs** tab. Expand the **Database** folder, then select **SAP HANA**.
2. Click **New**, then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select a template. Available options include **Instant Database Restore** and **Instant Disk Restore**.
5. Click **Source**. From the drop-down menu select **Application Browse** to select a source site and an application server to view available database recovery points. Select resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **Application Search** from the drop-down menu to search for application servers with available recovery points. Add copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy**. Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Destination**. Select a source site and an associated destination. Review the destination's database name mapping settings.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
9. To edit options before creating the job definition, click **Advanced**. Set the following job definition options.

Application Options

- Rename Mount Points – For more information about the **Rename Mount Points** options, see [Restore Jobs - Rename Mount Points and Initialization Parameter Options](#).

Policy Options

- Continue with next source on failure – Toggle the recovery of a resource in a series if the previous resource recovery fails. If unselected, the Restore job stops if the recovery of a resource fails.
- Automatically clean up resources on failure – Enable to automatically clean up allocated resources as part of a restore if the database recovery fails.
- Allow to overwrite and force clean up of pending old sessions – Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.
- Allow to overwrite vDisk – In cases where the Make Permanent option is enabled, and the destination VM has conflicting VMDK files, enable the **Allow to overwrite vDisk** option to delete the existing VMDK and overwrite it with the selected source.
- Job-Level Scripts – Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#). **State** and **status** arguments are not supported for Backup jobs.

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

Storage Options

- Make Permanent – Set the default permanent restoration action of the job. All database recovery operations can leverage Instant or Test modes and then either be deleted or promoted to permanent mode. This behavior is controlled through the **Make Permanent** option.
 - Enabled – Always make permanent through full copy FlashCopy
 - Disabled – Never make permanent
 - User Selection – Allows the user to select Make Permanent or Cleanup when the job session is pending

Important:

- When you run a restore with **Make Permanent**, the data is vMotioned to the datastore where the virtual machine's (VM) VMX file resides. The datastore to which the data and logs are moved may not be ideal or even supported by IBM Storage Copy Data Management for future backups. In most cases, the data and logs will be moved to the same datastore as the VM operating system disk and this can result in subsequent backup failures. You should inspect the VM configuration after restore with **Make Permanent** completes and manually reconfigure the VM to move the data and logs disks to datastores that are supported for subsequent backups and not the datastore containing the VM operating system disk. Then run VM and application inventory jobs explicitly to capture the updated configuration for the application servers. Finally, you can run another backup job of the resource so that a snapshot is available for future restore jobs.
 - When you run a restore with Make Permanent, it is recommended that an entry is added in the `/etc/fstab` file for your later reference.
 - Revert – Set the source production SAP HANA machines to the snapshot. This may have an impact on database downtime and could present a risk to the source production machine.
 - Enabled – Always revert
 - Disabled – Never revert
 - User Selection – Allows the user to select Revert or Cleanup when the job session is completed
- Note:** IBM Storage Copy Data Management does not support the **Revert** function if the backup was created using an IBM Storage Virtualize for snapshot SLA policy. The **Revert** function is only available for backups that were created using the IBM Storage Virtualize SLA policy.
- Protocol Priority – If more than one storage network protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

What to do next

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

Related tasks

[“Creating a File System Backup job” on page 129](#)

The procedures describe how to create a File System Backup job definition.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Creating an Oracle Restore job definition

The procedures describe how to create an Oracle Restore job definition.

Before you begin

- Create and run an Oracle Backup job. See [Create a Backup Job Definition - Oracle](#).
- Review Oracle requirements. See [Oracle requirements](#).
- To ensure that filesystem permissions are retained correctly when IBM Storage Copy Data Management moves Oracle data between servers, ensure that the user and group IDs of the Oracle users (e.g. oracle, oinstall, dba) are consistent across all the servers. Refer to Oracle documentation for recommended uid and gid values.
- If Oracle data resides on LVM volumes, you must stop and disable the lvm2-lvmetad service before running Backup or Restore jobs. Leaving the service enabled can prevent volume groups from being resigatured correctly during restore and can lead to data corruption if the original volume group is also present on the same system. To disable the lvm2-lvmetad service, run the following commands:

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

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

```
use_lvmetad = 0
```

- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).
- You must add credentials to the destination virtual machine when recovering with the `subnet` option. See [Add Credentials to a Virtual Machine](#).

Oracle database considerations:

- For Oracle 12c databases, backups are created without placing the database in hot backup mode through Oracle Storage Snapshot Optimization. All associated snapshot functionality is supported. This feature requires the Advanced Compression feature of Oracle to be licensed. If this feature is not licensed in your environment, perform the following procedure to disable Snapshot Optimization and force the use of hot backup mode:

Create the file `/etc/guestapps.conf` on the Oracle server and add the following to it:

```
[DEFAULT]
skipSnapshotOptimization = true
```

If the file already exists, edit it and add the parameter under the existing `[DEFAULT]` section. This is a per-host setting. The parameter must be set in this file on each Oracle server where you want to force the use of hot backup mode.

- Note that point-in-time recovery is not supported when one or more datafiles are added to the database in the period between the chosen point-in-time and the time that the preceding Backup job ran.
- To properly recover Oracle databases, the `ORACLE_BASE` environment variable must be set in the login environment of the user that owns the Oracle home. In cases where multiple homes exist on a node, each with a different owner, the variable must be set for each of the owners.
- During an Instant Database Restore, there may be failures if the new name specified for the restored database is similar to an existing database only differing by numerical suffix. For clustered instances of Oracle databases, the appliance always uses global database name in the UI. During the inventory and restore processes, individual instances using the numerical suffixes of the cluster must be correlated to the global database name. The issue with this comes when, as an example, “Production12” is discovered. Is this the instance 12 of the “Production” database, or instance two (2) of the “Production1” database, or a database named “Production12.”

About this task

IBM Storage Copy Data Management leverages Copy Data Management technology for recovering application databases through Database Restore jobs. Your Oracle clones can be utilized and consumed instantly through IBM Storage Copy Data Management Instant Disk Restore jobs. IBM Storage Copy Data Management catalogs and tracks all cloned instances. Instant Disk Restore leverages ISCSI or fibre channel protocols to provide immediate mount of LUNs without transferring data. Snapshotted databases are cataloged and instantly recoverable with no physical transfer of data. Point-in-time recovery is supported with log forwarding.

The following Oracle Database workflows are supported: **DevOps**, which provides Instant Disk Restore or Instant Recovery to a new location using a masked image, created through a Restore job with data masking enabled, **Instant Database Recovery**, which provides Instant Disk Restore or Instant Recovery using a non-masked image and point-in-time transaction logs, and **Instant Disk Restore**, which mounts a database for RMAN restores with application and operating system support.

Tip: Create a schedule before creating a job definition so that you can easily add the schedule to the job definition.

Procedure

1. Click the **Jobs** tab. Expand the **Database** folder, then select **Oracle**.

2. Click **New**, then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select a template. Available options include **DevOps**, **Instant Database Restore** and **Instant Disk Restore**.
5. Click **Source**. From the drop-down menu select **Application Browse** to select a source site and an application server to view available database recovery points. Select resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **Application Search** from the drop-down menu to search for application servers with available recovery points. Add copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy**. Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.

If creating an **Instant Database Restore** job definition, an additional recovery option is available through the Select Version feature. Enable **Allow Point-in-Time selection when job runs** to leverage archived logs and enable a point-in-time recovery of the databases.

If creating an **Instant Disk Restore** job definition, the RMAN tag displays next to the time in the Version field. An Oracle administrator can correlate the RMAN backups to the IBM Storage Copy Data Management during job creation.

7. Click **Destination**. Select a source site and an associated Oracle home. Click the Destination field to enter an optional alternate name for the database.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Application Options

Record mounted copies in RMAN local repository

This option is available for Instant Disk Restore workflows.

Select this option to catalog mounted copies into RMAN at the end of the Instant Disk Restore job. This option is an alternative to the cataloging of IBM Storage Copy Data Management-created copies during a Backup job. If the RMAN cataloging option is selected in the Backup job definition, every copy created by IBM Storage Copy Data Management is cataloged in the source database immediately after the copy is created. By contrast, this option allows you to perform cataloging on-demand only for a specific copy and only when you intend to restore data from that copy through RMAN. Note that for the cataloging to succeed the target database must be running at the time the Instant Disk Restore job runs.

Rename Mount Points and Database Initialization Parameters

For more information about the **Rename Mount Points** and **Database Initialization Parameters** (Instant Database Recovery and DevOps workflows only) options, see [Restore Jobs - Rename Mount Points and Initialization Parameter Options](#).

ASM Disk Names

This option allows you to specify the disk naming pattern for restored ASM disks, if available. If **Use default pattern** is selected, IBM Storage Copy Data Management uses the default naming pattern, which is `/dev/ecx-asmdisk/*` in Linux environments or `/dev/ecx_asm*` in AIX environments. Select **Specify a custom pattern** to set ASM disks to follow any naming conventions that may be in use for existing disks. The custom pattern must begin with `"/dev"` and must end with an asterisk (`*`). During restore, IBM Storage Copy Data Management creates a device alias, or symlink, matching the specified pattern and replaces the asterisk with a unique disk name.

Tip: This option has no effect if the database being restored does not use any ASM disks.

Policy Options

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If unselected, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the database recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Replace existing database

This option is available for Instant Database Recovery workflows.

Select this option to replace an existing database with the same name during recovery. When an Instant Database Recovery is performed for a database and another database with the same name is already running on the destination host/cluster, IBM Storage Copy Data Management shuts down the existing database before starting up the recovered database. If this option is not selected, the Instant Database Recovery fails when IBM Storage Copy Data Management encounters an existing running database with the same name.

Leave database shut down after recovery

This option is available for Instant Database Recovery and DevOps workflows.

Select this option to shut down the recovered database once the recovery operation completes. The database can be started up manually once needed.

Allow to overwrite vDisk

In cases where the Make Permanent option is enabled, and the destination VM has conflicting VMDK files, enable the **Allow to overwrite vDisk** option to delete the existing VMDK and overwrite it with the selected source.

Do not stretch recovery disks (storage features like Enhance Stretched Cluster)

Enabling this option will not create stretched recovery disks. This only applies to volumes that have been created using stretched FlashCopies.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#). **State** and **status** arguments are not supported for Backup jobs.

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

Storage Options

Make Permanent

Set the default permanent restoration action of the job. All database recovery operations can leverage Instant or Test modes and then either be deleted or promoted to permanent mode. This behavior is controlled through the **Make Permanent** option.

Important: When you run a restore with **Make Permanent**, the data is vMotioned to the datastore where the virtual machine's (VM) VMX file resides. The datastore to which the data and logs are moved may not be ideal or even supported by IBM Storage Copy Data Management for future backups. In most cases, the data and logs will be moved to the same datastore as the VM operating system disk and this can result in subsequent backup failures. You should inspect the VM configuration after restore with **Make Permanent** completes and manually reconfigure the VM to move the data and logs disks to datastores that are supported for subsequent backups and not the datastore containing the VM operating system disk. Then run VM and application inventory jobs explicitly to capture the updated configuration for the application servers. Finally, you can run another backup job of the resource so that a snapshot is available for future restore jobs.

Enabled - Always make permanent through full copy FlashCopy

Disabled - Never make permanent

User Selection - Allows the user to select Make Permanent or Cleanup when the job session is pending

Protocol Priority

If more than one storage network protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

What to do next

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

Related tasks

[“Creating an Oracle Backup job definition” on page 116](#)

The procedures describe how to create an Oracle Backup job definition.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Creating a Microsoft SQL Restore job definition

The procedures describe how to create a Microsoft SQL Restore job definition.

Before you begin

- Create and run a SQL Backup job. See [Create a Backup Job Definition - SQL](#).
- Review SQL requirements. See [Microsoft SQL requirements](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).
- You must add credentials to the destination virtual machine when recovering with the subnet option. See [Add Credentials to a Virtual Machine](#).

Microsoft SQL server considerations:

- You must use the **Overwrite existing database** option to restore a Microsoft SQL database to a source instance. Note that in an Instant Disk Restore or Instant Database Restore job, the restore destination cannot be the source instance if this option is not selected. Due to Microsoft SQL restrictions, two databases with the same ID cannot exist on the same instance, including if the restored database is renamed through the job definition.
- You must use a proxy node when performing a restore to a SQL Failover cluster. Windows requires signatures to be unique. When a disk is attached that has a signature equal to one that is already attached, Windows keeps the disk in “offline” mode and does not read its partition table or mount its volumes. To prevent the issue of disk signature collision, during Instant Database Restore, IBM Storage Copy Data Management leverages a Windows proxy server to temporarily mount disks from snapshots, generate a new signature, then mount to original server.
- You must manually create a directory to store VSS provider logs when running IBM Storage Copy Data Management 2.2.6 and earlier. Create the following directory structure on the SQL server:
c:\temp\CDM\logs
- Note that Instant Disk Restores of AlwaysOn databases are restored to local instances only. Instant Database Recoveries of AlwaysOn databases are joined to destination Availability Groups if the restore target is a primary replica, and restored to the local instance in a "RESTORING" state if the restore target is a secondary replica.
- When running SQL Server 2012, certain Availability Group properties must be modified in order for log backups to generate in their specified location. Under **Backup Preferences**, set the **Where should backups occur?** option to **Any Replica**. Under **Replica backup priorities**, set the primary server instance to 55 and the secondary server instance to 45.
- It is recommended to keep the SQL transaction log location on a different disk than the SQL database data and log files.
- When creating an Instant Seeding restore job definition, the destination must be a non-system drive. The SQL database and log files must be located on non-system drives.

About this task

IBM Storage Copy Data Management leverages Copy Data Management technology for recovering application databases through Database Restore jobs. Your SQL clones can be utilized and consumed instantly through IBM Storage Copy Data Management Instant Disk Restore jobs. IBM Storage Copy Data Management catalogs and tracks all cloned instances. Instant Disk Restore leverages iSCSI or fibre channel protocols to provide immediate mount of LUNs without transferring data. Snapshotted databases are cataloged and instantly recoverable with no physical transfer of data. Point-in-time recovery is supported with log forwarding.

The following Microsoft SQL Server workflows are supported:

Instant Database Restore - Provides instant recovery using a non-masked image and point-in-time transaction logs.

When performing a restore to a primary replica of a specific AlwaysOn Availability Group, the restore database is added to the destination availability group, however, the secondary replica is not automatically replicated.

When performing a restore to a secondary replica, the database is restored with the "norecovery" option, and the database is set to a "restoring" state. If managing transaction log backups without using IBM Storage Copy Data Management, you can manually restore log files, and add the database to an availability group, assuming the lsn of the secondary and primary database copies meet the criteria.

Instant Disk Restore - Mounts a database for restores with application and operating system support. An Instant Disk Restore of an AlwaysOn database is restored to the local destination instance.

Instant Seeding - Instant Seeding is used when a primary database is up and running, a secondary replica exists (but no secondary database). The primary database can be protected through an IBM Storage Copy Data Management Copy job, with log backups performed regularly. The seeding process restores the database to the copy time and applies the transaction logs of the primary database. The database is then added to an availability group at the end of the restore.

Instant Seeding is a recovery template available for SQL AlwaysOn Availability Groups that initializes secondary replica instances using storage copies. During an Instant Seeding restore, snapshot data along with log backups are restored to the secondary database, then added to the AlwaysOn Availability Group. Other AlwaysOn Availability Group restore jobs place databases in "restoring" states or only create local instances. Additional processes at the end of the Instant Seeding restore job add the database of a secondary replica to the Availability Group.

When creating the Instant Seeding job definition, primary and secondary replicas of the original Availability Group are displayed as valid restore destinations. All log backups after the copy job time are restored; no point-in-time options are required. For more information, see the Instant Seeding Prerequisites and Limitations section of the [Microsoft SQL requirements](#).

Tip: Create a schedule before creating a job definition so that you can easily add the schedule to the job definition.

Procedure

1. Click the **Jobs** tab. Expand the **Database** folder, then select **SQL**.
2. Click **New**, then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select **Microsoft SQL (Standalone and Failover Cluster)** or **Microsoft SQL (Always On Availability Group)**.
5. Select a template. Available options include **Instant Database Restore** for Microsoft SQL Standalone/Failover Cluster and Always On Availability Group jobs, **Instant Disk Restore**, and **Instant Seeding** for Microsoft SQL Always On Availability Group jobs.
6. Click **Source**. From the drop-down menu select **Application Browse** to select a source site and an application server to view available database recovery points. Select resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **Application Search** from the drop-down menu to search for application servers with available recovery points. Add copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

7. Click **Copy**. Click on **Select Backup Date Range** to provide a date range for which recovery points are displayed. Select a range of dates by selecting a date for Start Time and a date for End Time. After a range is established, click **OK**. Sites containing copies of the selected data display. Select a site. By default, the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**.

Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.

If creating an **Instant Database Restore** job definition, an additional recovery option is available through the Select Version feature. Enable **Allow Point-in-Time selection when job runs** to leverage archived logs and enable a point-in-time recovery of the databases.

8. Click **Destination**. Select a source site and an associated Microsoft SQL database. Click the **New database name** field to enter an optional alternate name for the database. If the destination is a SQL Failover Cluster instance, select a Windows server proxy in the **Select a Windows server to resignature LUNs** section. When resignaturing a copy, IBM Storage Copy Data Management retains the data and mounts the volume to the proxy selected in the Select a Windows server to resignature LUNs section.

Tip:

- Any Windows node with iSCSI or Fibre Channel access to the storage can be selected as a proxy server, provided that the node is not part of the original cluster. It is recommended to select a standalone virtual or physical Windows node as a proxy server.
 - If creating an Instant Seeding job definition, the seeding target must be a secondary replica. If a primary AlwaysOn node is selected as a seeding target, the job will fail. The New database name option does not apply to Instant Seeding jobs.
9. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
 10. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Application Options

Roll back uncommitted transactions and leave the database ready to use

Select this option to restore the database to an online state. If selected, additional transaction logs cannot be restored. If deselected, uncommitted transactions are not rolled back, leaving the database non-operational. Additional transaction logs can then be restored.

Tip: This option does not apply to Instant Seeding jobs.

Overwrite existing database

Select this option to replace an existing database with the same name during recovery. When an Instant Database Recovery is performed for a database and another database with the same name is already running on the destination host/cluster, IBM Storage Copy Data Management shuts down the existing database before starting up the recovered database. If this option is not selected, the Instant Database Recovery fails when IBM Storage Copy Data Management encounters an existing running database with the same name.

Copy databases

Select this option to restore SQL files to the original file path that is currently in use. The IBM Storage Copy Data Management agent copies the database files from the snapshot or clone volumes to the database folders on the local drive.

Rename mount points

By default, IBM Storage Copy Data Management renames mount points to the SQL data directory of the target SQL instance. You can override this behavior through the Rename mount points option.

Default SQL data directory - Mount points are not renamed.

Original mount point or drive letter - The original volume mount point of the databases is used. This will keep the same drive mapping, but the database recovery will occur on an alternate server. If the original volume mount point cannot be used to mount a volume (for example, if the folder is not empty), the restore fails. Note that it is recommended to keep the SQL transaction log location on a different disk than the SQL database data and log files.

Add a custom mount point - If enabled, a custom prefix can be entered in the Prefix string field. The prefix must specify a valid, preexisting volume drive letter that can contain a volume mount point on the drive. The prefix substitutes the default root folder of the mount. For example, if the

default root folder of the mount drive is E:\SQLDataFiles\mnt, and F:\RestoreMnt is entered in the Mount Point Prefix field, E:\SQLDataFiles\mnt is renamed to F:\RestoreMnt. The drive, in this case the F drive, must exist on the destination server. If the folder, in this case the RestoreMnt folder, does not exist, it will be created.

Tip: This option does not apply to Instant Seeding jobs. Instant Seeding supports restoring to the original path and volume drive letter/mount points.

Policy Options

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If unselected, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the database recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Allow to overwrite vDisk

In cases where the Make Permanent option is enabled, and the destination VM has conflicting VMDK files, enable the **Allow to overwrite vDisk** option to delete the existing VMDK and overwrite it with the selected source.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#).

State and **status** arguments are not supported for Backup jobs.

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

Storage Options

Make Permanent

Set the default permanent restoration action of the job. All database recovery operations can leverage Instant or Test modes and then either be deleted or promoted to permanent mode. This behavior is controlled through the **Make Permanent** option.

Important: When you run a restore with **Make Permanent**, the data is vMotioned to the datastore where the virtual machine's (VM) VMX file resides. The datastore to which the data and logs are moved may not be ideal or even supported by IBM Storage Copy Data Management for future backups. In most cases, the data and logs will be moved to the same datastore as the VM operating system disk and this can result in subsequent backup failures. You should inspect the VM configuration after restore with **Make Permanent** completes and manually reconfigure the VM to move the data and logs disks to datastores that are supported for subsequent backups and not

the datastore containing the VM operating system disk. Then run VM and application inventory jobs explicitly to capture the updated configuration for the application servers. Finally, you can run another backup job of the resource so that a snapshot is available for future restore jobs.

Enabled - Always make permanent through full copy FlashCopy

Disabled - Never make permanent

User Selection - Allows the user to select Make Permanent or Cleanup when the job session is pending

Protocol Priority

If more than one storage network protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

11. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

12. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

13. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

What to do next

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

Related tasks

[“Creating a SQL Backup job definition” on page 125](#)

The procedures describe how to create a SQL Backup job definition.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Creating a File System Restore job definition

The procedures describe how to create a File System Restore job definition.

Before you begin

- Review File System requirements. See [File system requirements](#).

- Create and run a File System Backup job. See [Create a Backup Job Definition - File System](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).
- You must add credentials to the destination virtual machine when recovering with the subnet option. See [Add Credentials to a Virtual Machine](#).

About this task

IBM Storage Copy Data Management leverages Copy Data Management technology for recovering physical Windows, AIX, and Linux file systems through File System Restore jobs. Clones can be utilized and consumed instantly through IBM Storage Copy Data Management Instant Disk Restore jobs. IBM Storage Copy Data Management catalogs and tracks all cloned instances. Instant Disk Restore jobs mount file systems from storage copies containing application data.

Procedure

1. Click the **Jobs** tab. Expand the **File System** folder, then select **File System**.
2. Click **New**, then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Instant Disk Restore** template.
5. Click **Source**. From the drop-down menu select **Application Browse** to select a source site and an application server to view available database recovery points. Select resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **Application Search** from the drop-down menu to search for application servers with available recovery points. Add copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy**. Sites containing copies of the selected data display. Select a site. By default, the latest copy will be used. If you would like to choose a specific version or use latest successful scan, select a site then click the **Select Specific Version or Use Latest Successful Scan**. Click the **Use Latest** or **Use Latest Successful Scan** to view specific copies and their associated job and completion time. If a selected recovery of a snapshot fails, another selection for the recovery should be done manually in the restore job, which will be a copy from the same site being used.
7. Click **Destination**. Select a source site and an associated destination. Review the destination's mount point mapping settings. Optionally, click the **Enter an alternate mount point** field to create an alternate mount point, or select **Use original mount points**.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Application Options

Overwrite Existing Mount Points

Select to overwrite the mount points at their original location.

Policy Options

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If unselected, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the database recovery fails.

Allow to overwrite vDisk

In cases where the Make Permanent option is enabled, and the destination VM has conflicting VMDK files, enable the **Allow to overwrite vDisk** option to delete the existing VMDK and overwrite it with the selected source.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#). **State** and **status** arguments are not supported for Backup jobs.

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

Tip: If adding a script to a Windows-based File System job definition, the user running the script must have the "Log on as a service" right enabled, which is required for running prescripts and postscripts. For more information about the "Log on as a service" right, see [Add the Log on as a service Right to an Account](#).

Storage Options

Make Permanent

Set the default permanent restoration action of the job. All database recovery operations can leverage Instant or Test modes and then either be deleted or promoted to permanent mode. This behavior is controlled through the **Make Permanent** option.

Important: When you run a restore with **Make Permanent**, the data is vMotioned to the datastore where the virtual machine's (VM) VMX file resides. The datastore to which the data and logs are moved may not be ideal or even supported by IBM Storage Copy Data Management for future backups. In most cases, the data and logs will be moved to the same datastore as the VM operating system disk and this can result in subsequent backup failures. You should inspect the VM configuration after restore with **Make Permanent** completes and manually reconfigure the VM to move the data and logs disks to datastores that are supported for subsequent backups and not the datastore containing the VM operating system disk. Then run VM and application inventory jobs explicitly to capture the updated configuration for the application servers. Finally, you can run another backup job of the resource so that a snapshot is available for future restore jobs.

Enabled - Always make permanent through full copy FlashCopy

Disabled - Never make permanent

User Selection - Allows the user to select Make Permanent or Cleanup when the job session is pending

Protocol Priority

If more than one storage network protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

What to do next

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

Related tasks

[“Creating a File System Backup job” on page 129](#)

The procedures describe how to create a File System Backup job definition.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Creating an IBM Storage Virtualize Restore job definition

IBM Storage Copy Data Management leverages Copy Data Management technology for recovering IBM Storage Virtualize volumes through IBM Storage Virtualize Restore jobs.

Instant Disk Restore

Provides instant writable access to a volume. An IBM Storage 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 Storage Copy Data Management IBM Storage Virtualize Backup job. Volumes can be restored to their original location or a new volume in the same or different IBM storage system.

Creating an Instant Disk Restore IBM Storage Virtualize Restore job definition

The procedures describe how to create an Instant Disk Restore IBM Storage Virtualize Restore job definition.

Before you begin

- Create and run one or more IBM Storage Virtualize Backup jobs. See [Create a Backup Job Definition - IBM Storage Virtualize](#).

- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

- IBM providers utilize port 22 for communication with IBM Storage Copy Data Management.
- Note that to restore data to an original volume, you must first offline the target disk on the host prior to recovery. Once recovery completes, bring the target disk back online.
- Note that after restoring data to an alternate location you must map the host to the restore volume on the IBM storage system. Then rescan the disk on the host, and bring the disk online.
- In IBM storage environments, port grouping and IP partnerships are required to enable remote copy connections. See IBM's SAN Volume Controller and Storwize Family Native IP Replication Guide.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#).

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **IBM Storage Virtualize**.
2. Click **New**, then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Instant Disk Restore** template.
5. Click **Source**. From the drop-down menu select **Volume** to select a source site and an associated IBM source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy**. Sites containing copies of the selected data display. Select a site. By default, the latest copy will be used. If you would like to choose a specific version or use latest successful scan, select a site then click the **Select Specific Version or Use Latest Successful Scan**. Click the **Use Latest** or **Use Latest Successful Scan** to view specific copies and their associated job and completion time. If a selected recovery of a snapshot fails, another selection for the recovery should be done manually in the restore job, which will be a copy from the same site being used.

Tip: When selecting a specific version, data created through VMware Backup jobs that apply to the selected IBM resource display, as the same data is contained with the snapshot for VMware and non-VMware related data.

7. Click **Destination**. Select the IBM hosts that contain the iSCSI Qualified Name (IQN) or Fibre Channel WWPN of the application that you want to assign to.

Tip: The IBM hosts that are used during runtime may be different based on the initiator name.

8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Make IA clone resource permanent

Enable to turn the snapshot copy into a proper resource that will not be cleaned up after the Instant Disk Restore job completes.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Allow to overwrite and force clean up of pending old sessions

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.

Skip waiting for FC to finish copying during storage controller restore

Enable to skip waiting for the FlashCopy to finish copying during a storage controller restore job. This option is disabled by default.

Do not stretch recovery disks (storage features like Enhance Stretched Cluster)

Enabling this option will not create stretched recovery disks. This only applies to volumes that have been created using stretched FlashCopies.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#). **State** and **status** arguments are not supported for Backup jobs.

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

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

What to do next

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

Related concepts

[“Using state and status arguments in postscripts” on page 206](#)

Related tasks

[“Creating an IBM Storage Virtualize Backup job definition” on page 131](#)

The procedures describe how to create an IBM Storage Virtualize Backup job definition.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Related reference

[“Search and filter guidelines” on page 339](#)

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.

Creating a Restore Volume(s) IBM Storage Virtualize Restore job definition

The procedures describe how to create a Restore Volume(s) IBM Storage Virtualize Restore job definition.

Before you begin

- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

- IBM providers utilize port 22 for communication with IBM Storage Copy Data Management.
- Note that to restore data to an original volume, you must first offline the target disk on the host prior to recovery. Once recovery completes, bring the target disk back online.
- Note that after restoring data to an alternate location you must map the host to the restore volume on the IBM storage system. Then rescan the disk on the host, and bring the disk online.
- In IBM storage environments, port grouping and IP partnerships are required to enable remote copy connections. See IBM's SAN Volume Controller and Storwize Family Native IP Replication Guide.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#).

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **IBM Storage Virtualize**.
2. Click **New**, then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Restore Volume(s)** template.
5. Click **Source**. From the drop-down menu select **Volume** to select a source site and an associated IBM source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy**. Sites containing copies of the selected data display. Select a site. By default, the latest copy will be used. If you would like to choose a specific version or use latest successful scan, select a site then click the **Select Specific Version or Use Latest Successful Scan**. Click the **Use Latest** or **Use Latest Successful Scan** to view specific copies and their associated job and completion time. If

a selected recovery of a snapshot fails, another selection for the recovery should be done manually in the restore job, which will be a copy from the same site being used.

Tip: When selecting a specific version, data created through VMware Backup jobs that apply to the selected IBM resource display, as the same data is contained with the snapshot for VMware and non-VMware related data.

7. Click **Destination**. To restore to the original volume, select **Restore to original volume**, or select **Restore to alternative location** and select a volume and associated pool. If no pool is selected, the pool with the largest amount of space available is chosen by default.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Skip waiting for FC to finish copying during storage controller restore

Enable to skip waiting for the FlashCopy to finish copying during a storage controller restore job. This option is disabled by default.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#). **State** and **status** arguments are not supported for Backup jobs.

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

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

What to do next

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

Related concepts

[“Using state and status arguments in postscripts” on page 206](#)

Related tasks

[“Creating an IBM Storage Virtualize Backup job definition” on page 131](#)

The procedures describe how to create an IBM Storage Virtualize Backup job definition.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Related reference

[“Search and filter guidelines” on page 339](#)

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.

Creating a Safeguarded Copy Restore

This procedure describes how to create a Safeguarded Copy (SGC) restore.

Before you begin

- Create and run a Safeguarded Copy backup. For more information, see [Creating a Safeguarded Copy SLA policy](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).



Warning: The Safeguarded copies will be removed from storage after the retention period has expired as defined in storage SLA. IBM Storage Copy Data Management may still display snapshots whose retention period has expired until the next run of maintenance job. Attempting to restore a snapshot that has expired will fail.

Considerations:

- IBM storage providers utilize port 22 for communication with IBM Storage Copy Data Management.
- 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.
- After restoring data to an alternate location you must map the host to the restore volume on the IBM storage system. Next, rescan the disk on the host to bring the disk online.
- In IBM storage environments, port grouping and IP partnerships are required to enable remote copy connections. See IBM's SAN Volume Controller and Storwize Family Native IP Replication Guide.

- One or more schedules might also be associated with a job. Job sessions run based on the triggers that are defined in the schedule. For more information, see [Create a Schedule](#).

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **IBM Storage Virtualize**.
2. Click **New**, then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Restore Volume(s)** template.
5. Click **Source**. From the drop-down menu, select **Volume** to select a source site and an associated IBM source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy**. Sites containing copies of the selected data display. Select a site. By default, the latest copy will be used. If you would like to choose a specific version or use latest successful scan, select a site then click the **Select Specific Version or Use Latest Successful Scan**. Click the **Use Latest** or **Use Latest Successful Scan** to view specific copies and their associated job and completion time. If a selected recovery of a snapshot fails, another selection for the recovery should be done manually in the restore job, which will be a copy from the same site being used.

Tip: When selecting a specific version, data that is created through VMware Backup jobs that apply to the selected IBM resource display, as the same data is contained with the snapshot for VMware and non-VMware related data.

7. Click **Destination**. To restore to the original volume, select **Restore to original volume**, or select **Restore to alternative location** and select a volume and associated pool. If no pool is selected, the pool with the largest amount of space available is chosen by default.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Skip waiting for FC to finish copying during storage controller restore

Enable to skip waiting for the FlashCopy to finish copying during a storage controller restore job. This option is disabled by default.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

After completion, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note that additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials that are required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#). **State** and **status** arguments are not supported for Backup jobs.

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

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

What to do next

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

Creating an IBM Storage Virtualize for Snapshot Restore job definition

IBM Storage Copy Data Management leverages Copy Data Management technology for recovering for Snapshot volumes through IBM Storage Virtualize for Snapshot Restore jobs.

Note: IBM Storage Virtualize for Snapshot does not support **Restore volume** function.

Instant Disk Restore

IBM Storage Virtualize

Provides instant writable access to a volume. An IBM Storage Copy Data Management snapshot is mapped to a target server where it can be accessed, copied, or put immediately into production use as needed.

Creating an Instant Disk Restore IBM Storage Virtualize for snapshot restore job definition

The procedures describe how to create an Instant Disk Restore IBM Storage Virtualize for Snapshot restore job definition.

Before you begin

- Create and run one or more IBM Storage Virtualize for Snapshot backup jobs. See [Create a Backup Job Definition - IBM Storage Virtualize](#).

- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

- IBM providers utilize port 22 for communication with IBM Storage Copy Data Management.
- Note that to restore data to an original volume, you must first offline the target disk on the host prior to recovery. Once recovery completes, bring the target disk back online.
- Note that after restoring data to an alternate location you must map the host to the restore volume on the IBM storage system. Then rescan the disk on the host, and bring the disk online.
- In IBM storage environments, port grouping and IP partnerships are required to enable remote copy connections. See IBM's SAN Volume Controller and Storwize Family Native IP Replication Guide.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#).

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **IBM Storage Virtualize for Snapshot**.
2. Click **New**, then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Instant Disk Restore** template.
5. Click **Source**. From the drop-down menu select **Volume** to select a source site and an associated IBM source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy**. Sites containing copies of the selected data display. Select a site. By default, the latest copy will be used. If you would like to choose a specific version or use latest successful scan, select a site then click the **Select Specific Version or Use Latest Successful Scan**. Click the **Use Latest** or **Use Latest Successful Scan** to view specific copies and their associated job and completion time. If a selected recovery of a snapshot fails, another selection for the recovery should be done manually in the restore job, which will be a copy from the same site being used.

Tip: When selecting a specific version, data created through VMware Backup jobs that apply to the selected IBM resource display, as the same data is contained with the snapshot for VMware and non-VMware related data.

7. Click **Destination**. Select the IBM hosts that contain the iSCSI Qualified Name (IQN) or Fibre Channel WWPN of the application that you want to assign to.

Tip: The IBM hosts that are used during runtime may be different based on the initiator name.

8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Make IA clone resource permanent

Enable to turn the snapshot copy into a proper resource that will not be cleaned up after the Instant Disk Restore job completes.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Allow to overwrite and force clean up of pending old sessions

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.

Skip waiting for FC to finish copying during storage controller restore

Enable to skip waiting for the FlashCopy to finish copying during a storage controller restore job. This option is disabled by default.

Do not stretch recovery disks (storage features like Enhance Stretched Cluster)

Enabling this option will not create stretched recovery disks. This only applies to volumes that have been created using stretched FlashCopies.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#). **State** and **status** arguments are not supported for Backup jobs.

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

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

What to do next

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

Creating an NetAPP ONTAP Restore job definition

IBM Storage Copy Data Management leverages Copy Data Management technology for recovering NetApp ONTAP volumes and files through NetApp ONTAP Restore jobs.

Instant Disk Restore

Provides instant writable access to volume or LUN. An IBM Storage 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 Storage Copy Data Management NetApp ONTAP Backup job. Volumes can be restored to their original location or a new volume in the same or different NetApp ONTAP cluster or server.

Note that Restore Volume jobs are not available for NetApp ONTAP storage systems operating in 7-mode. 7-mode resources will not display in Source or Destination steps.

Restore File(s)

Recover files from a primary snapshot created through an IBM Storage Copy Data Management NetApp ONTAP Backup job. Files are restored to their original location.

Creating an Instant Disk Restore NetApp ONTAP Restore job definition

The procedures describe how to create an Instant Disk Restore NetApp ONTAP Restore job definition.

Before you begin

- Create and run one or more NetApp ONTAP Backup jobs. See [Create a Backup Job Definition - NetApp ONTAP](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

- Target volumes and datastores can be automatically expanded through the autogrow feature if supported by the underlying storage.
- Note that a volume restore through a NetApp ONTAP Restore job that includes NetApp ONTAP storage systems operating in 7-mode is not supported. These providers will not display during job creation.
- Note that NetApp ONTAP Backup jobs can only vault or mirror snapshots created through IBM Storage Copy Data Management jobs.
- Note that a file restore from a mirror location is not available for NetApp ONTAP storage systems operating in 7-mode.
- Restore Volume jobs are not available for NetApp ONTAP storage systems operating in 7-mode. 7-mode resources will not display in Source or Destination steps.
- Note that a file restore through a NetApp ONTAP Restore job can only utilize the alternate location feature if both the source and the destination are NetApp ONTAP storage systems running Command Data ONTAP 8.3.
- Note that NetApp ONTAP storage systems operating in 7-mode support file recovery from primary snapshots to their original locations. To restore files from a mirror source, create and run an Instant Disk Restore job with the mirror as a source, then mount the restored volume via CIFS or NFS. Files can then be copied to a new location.
- Note that the .snapshot folder must be visible on NFS shares in order to properly view and run Copy Data Management jobs on NetApp ONTAP storage systems running Data ONTAP in 7-Mode or Clustered Data ONTAP up to and including version 8.2. Confirm with your administrator that the .snapshot folder is not hidden in your NetApp ONTAP environment.
- NetApp ONTAP and VMware Restore jobs will fail if the iSCSI Initiator Group (iGroup) is not configured on the NetApp Clustered Data ONTAP 8.3 storage system target. The procedure only needs to be performed once. Previously created iGroups for earlier versions of NetApp Clustered Data ONTAP do not

need to be reconfigured for version 8.3. Note that there should only be one iGroup using the software iSCSI initiator. For more information, contact Technical Support.

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

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **NetApp ONTAP**.
2. Click **New**, then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Instant Disk Restore** template.
5. Click **Source**. From the drop-down menu select **Volume** to select a source site and an associated NetApp ONTAP source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy**. Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version field** to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.

Tip: When selecting a specific version, data created through VMware Backup jobs that apply to the selected NetApp ONTAP resource display, as the same data is contained with the snapshot for VMware and non-VMware related data.

7. Click **Destination**. Select your NFS and CIFS mapping options, including the Volume Name Prefix. If a prefix is defined, the resulting NFS path or CIFS share displays as follows:

NFS: "/" + "volumeNamePrefix" + "_" + "sourcevolumeName"

CIFS: "volumeNamePrefix" + "_" + "sourcevolumeName"

If a prefix is not defined, a unique naming convention is applied. Unique prefixes should be defined for jobs that run concurrently. If a job is run with a specified prefix and the resulting Instant Disk Restore volume is made permanent, you must rename the NFS/CIFS path for the permanent volume prior to running the same job again.

8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Make IA clone resource permanent

Enable to turn the snapshot copy into a proper resource that will not be cleaned up after the instant access job completes.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the NetApp ONTAP volume recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#). **State** and **status** arguments are not supported for Backup jobs.

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

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

What to do next

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

Related concepts

[“Using state and status arguments in postscripts” on page 206](#)

Related tasks

[“Creating a NetApp ONTAP Backup job definition” on page 140](#)

The procedures describe how to create a NetApp ONTAP Backup job definition.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Related reference

[“Search and filter guidelines” on page 339](#)

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.

Creating a Restore Volume(s) NetApp ONTAP Restore job definition

The procedures describe how to create a Restore Volume(s) NetApp ONTAP Restore job definition.

Before you begin

- Create and run one or more NetApp ONTAP Backup jobs. See [Create a Backup Job Definition - NetApp ONTAP](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

- Target volumes and datastores can be automatically expanded through the autogrow feature if supported by the underlying storage.
- Note that a volume restore through a NetApp ONTAP Restore job that includes NetApp ONTAP storage systems operating in 7-mode is not supported. These providers will not display during job creation.
- Note that NetApp ONTAP Backup jobs can only vault or mirror snapshots created through IBM Storage Copy Data Management jobs.
- Note that a file restore from a mirror location is not available for NetApp ONTAP storage systems operating in 7-mode.
- Restore Volume jobs are not available for NetApp ONTAP storage systems operating in 7-mode. 7-mode resources will not display in Source or Destination steps.
- Note that a file restore through a NetApp ONTAP Restore job can only utilize the alternate location feature if both the source and the destination are NetApp ONTAP storage systems running Command Data ONTAP 8.3.
- Note that NetApp ONTAP storage systems operating in 7-mode support file recovery from primary snapshots to their original locations. To restore files from a mirror source, create and run an Instant Disk Restore job with the mirror as a source, then mount the restored volume via CIFS or NFS. Files can then be copied to a new location.
- Note that the .snapshot folder must be visible on NFS shares in order to properly view and run Copy Data Management jobs on NetApp ONTAP storage systems running Data ONTAP in 7-Mode or Clustered Data ONTAP up to and including version 8.2. Confirm with your administrator that the .snapshot folder is not hidden in your NetApp ONTAP environment.
- NetApp ONTAP and VMware Restore jobs will fail if the iSCSI Initiator Group (iGroup) is not configured on the NetApp Clustered Data ONTAP 8.3 storage system target. The procedure only needs to be performed once. Previously created iGroups for earlier versions of NetApp Clustered Data ONTAP do not need to be reconfigured for version 8.3. Note that there should only be one iGroup using the software iSCSI initiator. For more information, contact Technical Support.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#).

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **NetApp ONTAP**.

2. Click **New**, then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Restore Volume(s)** template.
5. Click **Source**. From the drop-down menu select **Volume** to select a source site and an associated NetApp ONTAP source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy**. Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version field** to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Destination**. To restore to the original volume, select **Restore to original volume**, or select **Restore to new volume in the same or different NetApp ONTAP cluster or server** and select a volume and associated aggregate. If no aggregate is selected, the aggregate with the largest amount of space available is chosen by default.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Auto mount NFS after volume restored

Enable to automatically mount the restored volume after restoration completes.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#). **State** and **status** arguments are not supported for Backup jobs.

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

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

What to do next

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

Related concepts

[“Using state and status arguments in postscripts” on page 206](#)

Related tasks

[“Creating a NetApp ONTAP Backup job definition” on page 140](#)

The procedures describe how to create a NetApp ONTAP Backup job definition.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Related reference

[“Search and filter guidelines” on page 339](#)

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.

Creating a Restore File(s) NetApp ONTAP Restore job definition

The procedures describe how to create a Restore File(s) NetApp ONTAP Restore job definition.

Before you begin

- Create and run one or more NetApp ONTAP Backup jobs. See [Create a Backup Job Definition - NetApp ONTAP](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

- Target volumes and datastores can be automatically expanded through the autogrow feature if supported by the underlying storage.

- Note that a volume restore through a NetApp ONTAP Restore job that includes NetApp ONTAP storage systems operating in 7-mode is not supported. These providers will not display during job creation.
- Note that NetApp ONTAP Backup jobs can only vault or mirror snapshots created through IBM Storage Copy Data Management jobs.
- Note that a file restore from a mirror location is not available for NetApp ONTAP storage systems operating in 7-mode.
- Restore Volume jobs are not available for NetApp ONTAP storage systems operating in 7-mode. 7-mode resources will not display in Source or Destination steps.
- Note that a file restore through a NetApp ONTAPRestore job can only utilize the alternate location feature if both the source and the destination are NetApp ONTAP storage systems running Command Data ONTAP 8.3.
- Note that NetApp ONTAP storage systems operating in 7-mode support file recovery from primary snapshots to their original locations. To restore files from a mirror source, create and run an Instant Disk Restore job with the mirror as a source, then mount the restored volume via CIFS or NFS. Files can then be copied to a new location.
- Note that the .snapshot folder must be visible on NFS shares in order to properly view and run Copy Data Management jobs on NetApp ONTAP storage systems running Data ONTAP in 7-Mode or Clustered Data ONTAP up to and including version 8.2. Confirm with your administrator that the .snapshot folder is not hidden in your NetApp ONTAP environment.
- NetApp ONTAP and VMware Restore jobs will fail if the iSCSI Initiator Group (iGroup) is not configured on the NetApp Clustered Data ONTAP 8.3 storage system target. The procedure only needs to be performed once. Previously created iGroups for earlier versions of NetApp Clustered Data ONTAP do not need to be reconfigured for version 8.3. Note that there should only be one iGroup using the software iSCSI initiator. For more information, contact Technical Support.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#).

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **NetApp ONTAP**.
2. Click **New**, then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Restore File(s)** template.
5. 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.
6. Click **Copy**. Sites containing copies of the selected files display. Select a site. The latest copy of the file is used. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Destination**. To restore to the original volume, select **Restore to original volume**, or select **Restore to alternate location** and select an alternate volume or directory.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#). **State** and **status** arguments are not supported for Backup jobs.

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

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

What to do next

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

Related concepts

[“Using state and status arguments in postscripts” on page 206](#)

Related tasks

[“Creating a NetApp ONTAP Backup job definition” on page 140](#)

The procedures describe how to create a NetApp ONTAP Backup job definition.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Related reference

[“Search and filter guidelines” on page 339](#)

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.

Creating a Pure Storage FlashArray Restore job definition

IBM Storage Copy Data Management leverages Copy Data Management technology for recovering Pure Storage FlashArray volumes through Pure Storage FlashArray Restore jobs.

Instant Disk Restore

Provides instant writable access to a volume. An IBM Storage 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 Storage Copy Data Management Pure Storage FlashArray Backup job. Volumes can be restored to their original location or a new volume in the same or different Pure Storage system.

Pure Storage FlashArray can be set to utilize CloudSnap functionality so that snapshots are offloaded from the local storage array to a S3 cloud storage target or NFS share. When a recovery job runs, the snapshot will be recovered from the local storage array. If it has been condensed out of the local storage and the **Add Snapshot Offload** option was set during SLA creation, the offload copy from cloud storage will be restored.

Tip: The IBM Storage Copy Data Management user interface will indicate that a Pure Storage FlashArray CloudSnap offload job has completed even though the transfer is still occurring in the background which is dependent on network speeds. Consider setting an age as the retention for offload copies when using the Pure Storage FlashArray CloudSnap functionality. Doing so will ensure that a sufficient amount of time has passed for data to be transferred to the S3 storage target or NFS share before it has to be condensed out from a backup. This is particularly important if several offload jobs are run in quick succession.

Creating an Instant Disk Restore Pure Storage FlashArray Restore job definition

The procedures describe how to create an Instant Disk Restore Pure Storage FlashArray Restore job definition.

Before you begin

- Create and run one or more Pure Storage FlashArray Backup job. See [Create a Backup Job Definition - Pure Storage FlashArray](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

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

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **Pure Storage FlashArray**.
2. Click **New**, then select **Restore**. The job editor opens.

3. Enter a name for your job definition and a meaningful description.
4. Select the **Instant Disk Restore** template.
5. Click **Source**. From the drop-down menu select **Volume** to select a source site and an associated Pure Storage source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy**. Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version field** to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Destination**. Select a Pure Storage FlashArray destination.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#). **State** and **status** arguments are not supported for Backup jobs.

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

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

What to do next

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

Related concepts

[“Using state and status arguments in postscripts” on page 206](#)

Related tasks

[“Creating a Pure Storage FlashArray Backup job definition” on page 142](#)

The procedures describe how to create a Pure Storage FlashArray Backup job definition.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Related reference

[“Search and filter guidelines” on page 339](#)

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.

Creating a Restore Volume(s) Pure Storage FlashArray Restore job definition

The procedures describe how to create a Restore Volume(s) Pure Storage FlashArray Restore job definition.

Before you begin

- Create and run one or more Pure Storage FlashArray Backup job. See [Create a Backup Job Definition - Pure Storage FlashArray](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).

Considerations:

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

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **Pure Storage FlashArray**.
2. Click **New**, then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Restore Volume(s)** template.
5. Click **Source**. From the drop-down menu select **Volume** to select a source site and an associated Pure Storage source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy**. Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version field** to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Destination**. To restore to the original volume, select **Restore to original volume**, or select **Restore to alternative location** and select a volume and associated pool. If no pool is selected, the pool with the largest amount of space available is chosen by default.
8. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
9. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the volume recovery fails.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#). **State** and **status** arguments are not supported for Backup jobs.

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

10. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

11. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

12. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

What to do next

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

Related concepts

[“Using state and status arguments in postscripts” on page 206](#)

Related tasks

[“Creating a Pure Storage FlashArray Backup job definition” on page 142](#)

The procedures describe how to create a Pure Storage FlashArray Backup job definition.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Related reference

[“Search and filter guidelines” on page 339](#)

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.

Creating a VMware Restore job definition

IBM Storage Copy Data Management leverages Copy Data Management technology for testing and cloning use cases, instant recovery, and full disaster recovery. VMware Backup jobs support Instant VM Restore and Instant Disk Restore scenarios.

Instant VM Restore jobs are run in the following modes:

Test Mode

Creates temporary virtual machines for development/testing, snapshot verification, and disaster recovery verification on a scheduled, repeatable basis without affecting production environments. Test machines are kept running as long as needed to complete testing and verification and are then cleaned up after testing and verification completes. Through fenced networking, you can establish a safe environment to test your jobs without interfering with virtual machines used for production. Virtual machines created through Test mode are also given unique names and identifiers to avoid conflicts within your production environment.

Clone Mode

Creates copies of virtual machines for use cases requiring permanent or long-running copies for data mining or duplication of a test environment in a fenced network. Virtual machines created through Clone mode are also given unique names and identifiers to avoid conflicts within your production environment. With clone mode you must be sensitive to resource consumption, since clone mode creates permanent or long-term virtual machines.

Production Mode

Enables disaster recovery at the local site from primary storage or a remote disaster recovery site, replacing original machine images with recover images. All configurations are carried over as part of the recovery, including names and identifiers, and all copy data jobs associated with the virtual machine continue to run.

You can also set an IP address or subnet mask for virtual machines to be repurposed for development/testing or disaster recovery use cases. Supported mapping types include IP to IP, IP to DHCP, and subnet to subnet.

Instant Disk Restore

Provides instant writable access to data and application recovery points. An IBM Storage Copy Data Management snapshot is mapped to a target server where it can be accessed, copied, or put immediately into production use as needed.

Tip: Create a schedule before creating a job definition so that you can easily add the schedule to the job definition.

Creating an Instant Disk Restore VMware Restore job definition

The procedures describe how to create an Instant Disk Restore VMware Restore job definition.

Before you begin

- Create and run a VMware Backup job. See [Create a Backup Job Definition - VMware](#).
- Ensure the latest version of VMware Tools is installed in your environment. IBM Storage Copy Data Management was tested against VMware Tools 9.10.0.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).
- You must add credentials to the destination virtual machine when recovering with the subnet option. See [Add Credentials to a Virtual Machine](#).
- To run network prescripts and postscripts, select the **Power on after recovery option**, then the **Enable pre and post network VM level scripts option**. Note that prior to running the job, the scripts must be defined and copied to specific locations on the virtual machine. In a Windows environment, copy `prenetwork.bat` and `postnetwork.bat` to your `c:\program files\ibm\IBM Storage Copy Data Management\scripts` directory. In a Linux environment, copy `prenetwork.sh` and `postnetwork.sh` to your `/opt/CDM/scripts/` directory.

Considerations:

- Note that VMware Backup and Restore jobs only support vCenters or ESX hosts running vSphere 6.0 through 7.0.
- When running VADP-based VM Replication workflows, target volumes and datastores can be automatically expanded in response to space usage requirements if supported by the underlying storage. Automatic growing prevents a volume from running out of space or forcing you to delete files manually. For a list of supported storage systems see [System requirements](#).
- Note that VMware DRS cluster datastores are supported in VMware Backup and Restore jobs.
- Note that after an Instant Disk Restore Restore job completes, your vDisk will be mounted but you may need to bring it online through the operating system from the Disk Management console.
- In addition to NFS, IBM Storage Copy Data Management supports VMFS datastores for NetApp ONTAP storage targets.

- Note that Instant Disk Restore recoveries utilizing the VM Replication method are not supported at the datastore level. Instant Disk Restore datastore level recoveries are supported through the primary storage snapshot method.
- Note that in Instant VM Restore recoveries utilizing NetApp ONTAP storage systems running Clustered Data ONTAP or Data ONTAP operating in 7-Mode, if a source with a swap directory on a dedicated datastore is recovered to a different destination, then the source datastore must have more free space than the amount of memory configured for the virtual machine. This may not be applicable if the virtual machine is configured with memory reservation.
- A VMware Restore job recovering a virtual machine from an ESX cluster protected with snapshot, vault, or mirror displays a Locate LUN failure if the maximum allowed LUNs for the ESX host recovery target reaches its limit.
- In NetApp ONTAP environments running Clustered Data ONTAP, cluster peering must be enabled. Peer relationships enable communication between SVMs. See NetApp ONTAP's Cluster and Vserver Peering Express Guide.
- In IBM storage environments, port grouping and IP partnerships are required to enable remote copy connections. See IBM's SAN Volume Controller and Storwize Family Native IP Replication Guide.
- NetApp ONTAP and VMware Restore jobs will fail if the iSCSI Initiator Group (iGroup) is not configured on the NetApp Clustered Data ONTAP 8.3 storage system target. The procedure only needs to be performed once. Previously created iGroups for earlier versions of NetApp Clustered Data ONTAP do not need to be reconfigured for version 8.3. For more information, contact Technical Support.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#).

Considerations for VMware volumes:

- All functionality of Restore workflows are supported where the original virtual machine was stored on a virtual volume (VVOL) datastore. Virtual machines can be recovered to a VVOL target or VMware datastore through Production or Clone mode.
- With the introduction of VVOLs, a storage vendor's vSphere API for Storage Awareness (VASA) stores metadata required to run virtual machines on VVOL datastores. IBM Storage Copy Data Management does not require the VASA provider metadata for Restore workflows, as IBM Storage Copy Data Management stores the VM Replication in a VMware target datastore. In case of a disaster in which the VASA provider is lost, a new VASA provider can be brought up to recover the virtual machine back to a VVOL datastore through Production or Clone mode.

Procedure

1. Click the **Jobs** tab. Expand the **Hypervisor** folder, then select **VMware**.
2. Click **New**, then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Instant Disk Restore** template.
5. Click **Source**, then select **VM Storage** or **Datastores** as the source type. Select a source site and an associated VMware source to view virtual machines, VM templates, folders, vApps, and datacenters with available recovery points. Select resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.
6. Click **Copy**. Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Destination**. Expand a VMware source to view virtual machines, folders, vApps, and datacenters available as destinations. To restore to the original host or cluster, select **Use original host or cluster**.
8. Select the datastore and virtual disk mapping options if you selected a destination different from the original host or cluster.

Virtual Disks

In the **VM** field, select virtual machine destinations.

In the **Disk Mode** field, select “persistent”, “independent persistent”, or “independent non-persistent”.

The values are defined as follows:

- **Persistent** - Changes are permanently written to the virtual disk. The disk is included in snapshots taken of its virtual machine.
- **Independent Persistent** - Changes are permanently written to the virtual disk. The disk is excluded from any snapshots taken of its virtual machine.
- **Independent Non-Persistent** - Changes to the virtual disk are discarded when the virtual machine powers off; the VMDK files revert to their original state. The disk is excluded from any snapshots taken of its virtual machine.

In the optional **Controller Type** field, select a supported SCSI controller, including LSI SAS, LSI Parallel, BusLogic, and VMware Paravirtual. Changing the SCSI controller type replaces the existing controller with a new controller, applies the common settings of the existing controller to the new controller, and reassigns all SCSI devices to the new controller.

Use the optional **Controller Address #** and **Controller LUN #** fields to select specific controllers or LUNs.

Datstores

Set the destination datastore.

9. To create the job definition using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the **Jobs** tab.
10. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Protocol Priority

If more than one storage protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

Make IA clone resource permanent

Enable to turn the snapshot copy into a proper resource that will not be cleaned up after the instant access job completes.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the virtual machine recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Do not stretch recovery disks (storage features like Enhance Stretched Cluster)

Enabling this option will not create stretched recovery disks. This only applies to volumes that have been created using stretched FlashCopies.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added

to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#). State and status arguments are not supported for Backup jobs.

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

11. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

12. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

13. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

What to do next

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

Related concepts

[“Using state and status arguments in postscripts” on page 206](#)

Related tasks

[“Creating a VMware Backup job definition” on page 145](#)

The procedures describe how to create a VMware Backup job definition.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Related reference

[“Search and filter guidelines” on page 339](#)

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.

Related information

Creating an Instant VM Restore VMware Restore job definition

The procedures describe how to create an Instant VM Restore VMware Restore job definition.

Before you begin

- Create and run a VMware Backup job. See [Create a Backup Job Definition - VMware](#).
- Ensure the latest version of VMware Tools is installed in your environment. IBM Storage Copy Data Management was tested against VMware Tools 9.10.0.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).
- You must add credentials to the destination virtual machine when recovering with the subnet option. See [Add Credentials to a Virtual Machine](#).
- To run network prescripts and postscripts, select the **Power on after recovery option**, then the **Enable pre and post network VM level scripts option**. Note that prior to running the job, the scripts must be defined and copied to specific locations on the virtual machine. In a Windows environment, copy `prenetwork.bat` and `postnetwork.bat` to your `c:\program files\ibm\IBM Storage Copy Data Management\scripts` directory. In a Linux environment, copy `prenetwork.sh` and `postnetwork.sh` to your `/opt/CDM/scripts/` directory.

Considerations:

- Note that VMware Backup and Restore jobs only support vCenters or ESX hosts running vSphere 6.0 through 7.0.
- When running VADP-based VM Replication workflows, target volumes and datastores can be automatically expanded in response to space usage requirements if supported by the underlying storage. Automatic growing prevents a volume from running out of space or forcing you to delete files manually. For a list of supported storage systems see [System requirements](#).
- Note that VMware DRS cluster datastores are supported in VMware Backup and Restore jobs.
- Note that after an Instant Disk Restore Restore job completes, your vDisk will be mounted but you may need to bring it online through the operating system from the Disk Management console.
- In addition to NFS, IBM Storage Copy Data Management supports VMFS datastores for NetApp ONTAP storage targets.
- Note that Instant Disk Restore recoveries utilizing the VM Replication method are not supported at the datastore level. Instant Disk Restore datastore level recoveries are supported through the primary storage snapshot method.
- Note that in Instant VM Restore recoveries utilizing NetApp ONTAP storage systems running Clustered Data ONTAP or Data ONTAP operating in 7-Mode, if a source with a swap directory on a dedicated datastore is recovered to a different destination, then the source datastore must have more free space than the amount of memory configured for the virtual machine. This may not be applicable if the virtual machine is configured with memory reservation.
- A VMware Restore job recovering a virtual machine from an ESX cluster protected with snapshot, vault, or mirror displays a Locate LUN failure if the maximum allowed LUNs for the ESX host recovery target reaches its limit.
- In NetApp ONTAP environments running Clustered Data ONTAP, cluster peering must be enabled. Peer relationships enable communication between SVMs. See NetApp ONTAP's Cluster and Vserver Peering Express Guide.

- In IBM storage environments, port grouping and IP partnerships are required to enable remote copy connections. See IBM's SAN Volume Controller and Storwize Family Native IP Replication Guide.
- NetApp ONTAP and VMware Restore jobs will fail if the iSCSI Initiator Group (iGroup) is not configured on the NetApp Clustered Data ONTAP 8.3 storage system target. The procedure only needs to be performed once. Previously created iGroups for earlier versions of NetApp Clustered Data ONTAP do not need to be reconfigured for version 8.3. For more information, contact Technical Support.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#).

Considerations for VMware volumes:

- All functionality of Restore workflows are supported where the original virtual machine was stored on a virtual volume (VVOL) datastore. Virtual machines can be recovered to a VVOL target or VMware datastore through Production or Clone mode.
- With the introduction of VVOLs, a storage vendor's vSphere API for Storage Awareness (VASA) stores metadata required to run virtual machines on VVOL datastores. IBM Storage Copy Data Management does not require the VASA provider metadata for Restore workflows, as IBM Storage Copy Data Management stores the VM Replication in a VMware target datastore. In case of a disaster in which the VASA provider is lost, a new VASA provider can be brought up to recover the virtual machine back to a VVOL datastore through Production or Clone mode.

Procedure

1. Click the **Jobs** tab. Expand the **Hypervisor** folder, then select **VMware**.
2. Click **New**, then select **Restore**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Select the **Instant VM Restore** template.
5. Click **Source**. From the drop-down menu, select **VMs and Templates** to choose a source site and an associated VMware source to view virtual machines, VM templates, datastores, folders, and vApps and with available recovery points. Select resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **VM Search** from the drop-down menu to search for virtual machines with available recovery points across all datacenters. Add virtual machine copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy**. Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Destination**. Select a destination site and an associated VMware destination to view virtual machines, folders, vApps, and datacenters available as destinations. To restore to the original host or cluster and allow IBM Storage Copy Data Management to define the destination IP address, select **Use original host or cluster with system defined IP configuration**. To restore to the original host or cluster using your predefined IP address configuration, select **Use original host or cluster with original IP configuration**. To restore to a destination different from the original host or cluster, select **Use alternative host or cluster**.
8. Select virtual network and datastore mapping options if you selected **Use alternative original host or cluster** in the previous step. The Virtual Networks pane displays all of the virtual networks associated with your VMware Restore job sources. New virtual networks must be selected for use at the recovery site, as well as new datastores on the Datastores pane. Select a production and test network in the Virtual Networks tab, and a destination datastore in the Datastore tab.

Virtual Networks

Set virtual networks for production and test recovery jobs. Destination network settings for production and test environments should be different locations.

Datastores

Set the destination datastore.

Subnet

Set an IP address or subnet mask for virtual machines to be repurposed for development/testing or disaster recovery use cases. Supported mapping types include IP to IP, IP to DHCP, and subnet to subnet. Virtual machines containing multiple NICs are supported.

By default, the **Use system defined subnets and IP addresses for VM guest OS on destination** option is enabled. To use your predefined subnets and IP addresses, select **Use original subnets and IP addresses for VM guest OS on destination**.

To create a new mapping configuration, select **Add mappings for subnets and IP addresses for VM guest OS on destination**, then click **Add Mapping**. Enter a subnet or IP address in the Source field. In the destination field, select **DHCP** to automatically select an IP and related configuration information if DHCP is available on the selected client. Select **Static** to enter a specific subnet or IP address, subnet mask, gateway, and DNS. Note that **Subnet** or **IP Address**, **Subnet Mask**, and **Gateway** are required fields. If a subnet is entered as a source, a subnet must also be entered as a destination.

IP reconfiguration is skipped for virtual machines if a static IP is used but no suitable subnet mapping is found, or if the source machine is powered off and there is more than one associated NIC. In a Windows environment, if a virtual machine is DHCP only, then IP reconfiguration is skipped for that virtual machine. In a Linux environment all addresses are assumed to be static, and only IP mapping will be available.

Tip: You must add credentials to the destination virtual machine when recovering with the subnet option. Note that if using a domain user account, the credentials must be added to the destination virtual machine, then configured through the **Test & Configure** option. This option verifies communication with the server, tests DNS settings between the IBM Storage Copy Data Management appliance and the server, and installs an IBM Storage Copy Data Management agent on the server. From the **Configure** tab in the Provider Browser pane, right-click the virtual machine, then click **Test & Configure**. See [Add Credentials to a Virtual Machine](#) and [Register a Provider](#).

VM Folder

Set the VM folder path on the destination datacenter. To use an alternate path, deselect **Use original VM folder path on destination** and enter the path in the **Use alternative path** field. Note that the directory will be created if it does not exist. Use "/" as the root VM folder of the targeted datacenter.

9. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
10. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Default Mode

Set the VMware Restore job to run in Test, Production, or Clone mode by default. Once the job is created, it can be run in Test, Production, or Clone mode through the Jobs tab.

Protocol Priority

If more than one storage protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

Power on after recovery

Toggle the power state of a virtual machine after a recovery is performed. Virtual machines are powered on in the order they are recovered, as set in the Source step. Turning this feature on also gives access to the Enable pre and post network VM level scripts option. Note that restored VM templates cannot be powered on after recovery.

Enable pre and post network VM level scripts

If the Power on after recovery option is enabled, network prescripts and postscripts can be run at the virtual machine level. Note that prior to running a VMware Backup job, the scripts must be defined and copied to specific locations on the virtual machine. In a Windows environment, copy prenetwork.bat and postnetwork.bat to your c:\program files\ibm\IBM Storage Copy Data

Management\scripts directory. In a Linux environment, copy prenetwork.sh and postnetwork.sh to your /opt/CDM/scripts/ directory.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the virtual machine recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field to add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#). State and status arguments are not supported for Backup jobs.

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

11. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

12. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

13. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.
14. Once the job completes successfully, select one of the following options from the **Actions** menu on the General tab of the job session on the Jobs tab: **End IV (Cleanup)**, **RRP (vMotion)**, or **Clone (vMotion)**.

End IV (Cleanup) destroys the virtual machine and cleans up all associated resources. Since this is a temporary/testing virtual machine, all data is lost when the virtual machine is destroyed.

RRP (vMotion) is equivalent to using the Production selection in the job Advanced screen. This option migrates the virtual machine through vMotion to the Datastore and the Virtual Network defined as the "For Production" Network.

Clone (vMotion) is equivalent to using the Clone selection in the job Advanced screen. This option migrates the virtual machine through vMotion to the Datastore and Virtual Network defined as the "For Test" network.

What to do next

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

Related concepts

[“Using state and status arguments in postscripts” on page 206](#)

Related tasks

[“Creating a VMware Backup job definition” on page 145](#)

The procedures describe how to create a VMware Backup job definition.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Related reference

[“Search and filter guidelines” on page 339](#)

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.

Related information

Creating an Instant VM Restore (Long Distance) VMware Restore job definition

The procedures describe how to create an Instant VM Restore (Long Distance) VMware Restore job definition.

Before you begin

- Create and run a VMware Backup job. See [Create a Backup Job Definition - VMware](#).
- Ensure the latest version of VMware Tools is installed in your environment. IBM Storage Copy Data Management was tested against VMware Tools 9.10.0.
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).
- You must add credentials to the destination virtual machine when recovering with the subnet option. See [Add Credentials to a Virtual Machine](#).
- To run network prescripts and postscripts, select the **Power on after recovery option**, then the **Enable pre and post network VM level scripts option**. Note that prior to running the job, the scripts must be defined and copied to specific locations on the virtual machine. In a Windows environment,

copy prenetwork.bat and postnetwork.bat to your c:\program files\ibm\IBM Storage Copy Data Management\scripts directory. In a Linux environment, copy prenetwork.sh and postnetwork.sh to your /opt/CDM/scripts/ directory.

Considerations:

- Note that VMware Backup and Restore jobs only support vCenters or ESX hosts running vSphere 6.0 through 7.0.
- When running VADP-based VM Replication workflows, target volumes and datastores can be automatically expanded in response to space usage requirements if supported by the underlying storage. Automatic growing prevents a volume from running out of space or forcing you to delete files manually. For a list of supported storage systems see [System requirements](#).
- Note that VMware DRS cluster datastores are supported in VMware Backup and Restore jobs.
- Note that after an Instant Disk Restore Restore job completes, your vDisk will be mounted but you may need to bring it online through the operating system from the Disk Management console.
- In addition to NFS, IBM Storage Copy Data Management supports VMFS datastores for NetApp ONTAP storage targets.
- Note that Instant Disk Restore recoveries utilizing the VM Replication method are not supported at the datastore level. Instant Disk Restore datastore level recoveries are supported through the primary storage snapshot method.
- Note that in Instant VM Restore recoveries utilizing NetApp ONTAP storage systems running Clustered Data ONTAP or Data ONTAP operating in 7-Mode, if a source with a swap directory on a dedicated datastore is recovered to a different destination, then the source datastore must have more free space than the amount of memory configured for the virtual machine. This may not be applicable if the virtual machine is configured with memory reservation.
- A VMware Restore job recovering a virtual machine from an ESX cluster protected with snapshot, vault, or mirror displays a Locate LUN failure if the maximum allowed LUNs for the ESX host recovery target reaches its limit.
- In NetApp ONTAP environments running Clustered Data ONTAP, cluster peering must be enabled. Peer relationships enable communication between SVMs. See NetApp ONTAP's Cluster and Vserver Peering Express Guide.
- In IBM storage environments, port grouping and IP partnerships are required to enable remote copy connections. See IBM's SAN Volume Controller and Storwize Family Native IP Replication Guide.
- NetApp ONTAP and VMware Restore jobs will fail if the iSCSI Initiator Group (iGroup) is not configured on the NetApp Clustered Data ONTAP 8.3 storage system target. The procedure only needs to be performed once. Previously created iGroups for earlier versions of NetApp Clustered Data ONTAP do not need to be reconfigured for version 8.3. For more information, contact Technical Support.
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#).

Considerations for VMware volumes:

- All functionality of Restore workflows are supported where the original virtual machine was stored on a virtual volume (VVOL) datastore. Virtual machines can be recovered to a VVOL target or VMware datastore through Production or Clone mode.
- With the introduction of VVOLS, a storage vendor's vSphere API for Storage Awareness (VASA) stores metadata required to run virtual machines on VVOL datastores. IBM Storage Copy Data Management does not require the VASA provider metadata for Restore workflows, as IBM Storage Copy Data Management stores the VM Replication in a VMware target datastore. In case of a disaster in which the VASA provider is lost, a new VASA provider can be brought up to recover the virtual machine back to a VVOL datastore through Production or Clone mode.

Procedure

1. Click the **Jobs** tab. Expand the **Hypervisor** folder, then select **VMware**.
2. Click **New**, then select **Restore**. The job editor opens.

3. Enter a name for your job definition and a meaningful description.
4. Select the **Instant VM Restore (Long Distance)** template.
5. Click **Source**. From the drop-down menu, select **VMs and Templates** to choose a source site and an associated VMware source to view virtual machines, VM templates, datastores, folders, and vApps and with available recovery points. Select resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid.

Alternatively, select **VM Search** from the drop-down menu to search for virtual machines with available recovery points across all datacenters. Add virtual machine copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.

6. Click **Copy**. Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version** field to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.
7. Click **Local Destination**. Select a local destination site and an associated VMware destination to view virtual machines, folders, vApps, and datacenters available as destinations. Note that the local storage destination is temporary, and is cleaned up after the job completes.
8. Click **Remote Destination**. Select a remote destination site and an associated VMware destination to view virtual machines, folders, vApps, and datacenters available as destinations.
9. Select virtual network and datastore mapping options. The Virtual Networks pane displays all of the virtual networks associated with your VMware Restore job sources. New virtual networks must be selected for use at the recovery site, as well as new datastores on the Datastores pane. Select a production and test network in the Virtual Networks tab, and a destination datastore in the Datastore tab.

Virtual Networks

Set virtual networks for production and test recovery jobs. Destination network settings for production and test environments should be different locations.

Datastores

Set the destination datastore.

Subnet

Set an IP address or subnet mask for virtual machines to be repurposed for development/testing or disaster recovery use cases. Supported mapping types include IP to IP, IP to DHCP, and subnet to subnet. Virtual machines containing multiple NICs are supported.

By default, the **Use system defined subnets and IP addresses for VM guest OS on destination** option is enabled. To use your predefined subnets and IP addresses, select **Use original subnets and IP addresses for VM guest OS on destination**.

To create a new mapping configuration, select **Add mappings for subnets and IP addresses for VM guest OS on destination**, then click **Add Mapping**. Enter a subnet or IP address in the Source field. In the destination field, select **DHCP** to automatically select an IP and related configuration information if DHCP is available on the selected client. Select **Static** to enter a specific subnet or IP address, subnet mask, gateway, and DNS. Note that **Subnet or IP Address**, **Subnet Mask**, and **Gateway** are required fields. If a subnet is entered as a source, a subnet must also be entered as a destination.

IP reconfiguration is skipped for virtual machines if a static IP is used but no suitable subnet mapping is found, or if the source machine is powered off and there is more than one associated NIC. In a Windows environment, if a virtual machine is DHCP only, then IP reconfiguration is skipped for that virtual machine. In a Linux environment all addresses are assumed to be static, and only IP mapping will be available.

Tip: You must add credentials to the destination virtual machine when recovering with the subnet option. Note that if using a domain user account, the credentials must be added to the destination virtual machine, then configured through the **Test & Configure** option. This option verifies communication with the server, tests DNS settings between the IBM Storage Copy Data Management appliance and the server, and installs an IBM Storage Copy Data Management

agent on the server. From the **Configure** tab in the Provider Browser pane, right-click the virtual machine, then click **Test & Configure**. See Add Credentials to a Virtual Machine on page 105 and Register a Provider on page 79.

VM Folder

Set the VM folder path on the destination datacenter. To use an alternate path, deselect **Use original VM folder path on destination** and enter the path in the **Use alternative path** field. Note that the directory will be created if it does not exist. Use "/" as the root VM folder of the targeted datacenter.

10. To create the job definition using default options, click **Create Job**. The job can be run manually from the **Jobs** tab.
11. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Default Mode

Set the VMware Restore job to run in Test, Production, or Clone mode by default. Once the job is created, it can be run in Test, Production, or Clone mode through the Jobs tab.

Protocol Priority

If more than one storage protocol is available, select the protocol to take priority in the job. Available protocols include iSCSI and Fibre Channel.

Power on after recovery

Toggle the power state of a virtual machine after a recovery is performed. Virtual machines are powered on in the order they are recovered, as set in the Source step. Turning this feature on also gives access to the **Enable pre and post network VM level scripts** option. Note that restored VM templates cannot be powered on after recovery.

Enable pre and post network VM level scripts

If the **Power on after recovery option** is enabled, network prescripts and postscripts can be run at the virtual machine level. Note that prior to running a VMware Backup job, the scripts must be defined and copied to specific locations on the virtual machine. In a Windows environment, copy prenetwork.bat and postnetwork.bat to your c:\program files\ibm\IBM Storage Copy Data Management\scripts directory. In a Linux environment, copy prenetwork.sh and postnetwork.sh to your /opt/CDM/scripts/ directory.

Continue with next source on failure

Toggle the recovery of a resource in a series if the previous resource recovery fails. If disabled, the Restore job stops if the recovery of a resource fails.

Automatically clean up resources on failure

Enable to automatically clean up allocated resources as part of a restore if the virtual machine recovery fails.

Allow to overwrite and force clean up of pending old sessions

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

Job-Level Scripts

Job-level pre-scripts and post-scripts are scripts that can be run before or after a job runs at the job-level. A script can consist of one or many commands, such as a shell script for Linux-based virtual machines or Batch and PowerShell scripts for Windows-based virtual machines.

In the Pre-Script and/or Post-Script section, click **Select** to select a previously uploaded script, or click **Upload** to upload a new script. Note that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure Scripts](#).

Once complete, the script displays in the Pre-Script or Post-Script section. Click the **Parameters** field at add a parameter to the script, then click **Add**. Note additional parameters can be added to a script by entering parameters one at a time in the field, then clicking **Add**. Next, click the **Identity** field to add or create the credentials required to run the script. Finally, click the **Application Server** field to define the location where the script will be injected and executed.

Repeat the above procedure to add additional Pre-Scripts and Post-Scripts. For information about script return codes, see [Return Code Reference](#).

For Restore job post-scripts only, the positional arguments **state** and **status** can be passed to the script. For information about this feature, see [Using State and Status Arguments in Postscripts](#). State and status arguments are not supported for Backup jobs.

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

12. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

13. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

14. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.
15. Once the job completes successfully, select one of the following options from the **Actions** menu on the General tab of the job session on the Jobs tab: **End IV (Cleanup)**, **RRP (vMotion)**, or **Clone (vMotion)**.

RRP (vMotion) is equivalent to using the Production selection in the job Advanced screen. This option migrates the virtual machine through vMotion to the Datastore and the Virtual Network defined as the "For Production" Network.

Clone (vMotion) is equivalent to using the Clone selection in the job Advanced screen. This option migrates the virtual machine through vMotion to the Datastore and Virtual Network defined as the "For Test" network.

What to do next

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

Related concepts

[“Using state and status arguments in postscripts” on page 206](#)

Related tasks

[“Creating a VMware Backup job definition” on page 145](#)

The procedures describe how to create a VMware Backup job definition.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Related reference

“Search and filter guidelines” on page 339

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.

Related information

Renaming mount points and initialization parameter options

This option is available for supported Restore jobs running Instant Database Restore and DevOps workflows.

Mount Point Rename

This option controls how mount points and ASM diskgroups are renamed when they are mounted on the destination host/cluster during recovery.

Append a timestamp

By default, IBM Storage Copy Data Management appends a timestamp to the original mount point. For example:

<i>Table 10. Mount Point Rename - Append a timestamp</i>	
Original Mount Point	New Mount Point
/u02/myproddb/data	/u02/myproddb/data_1479400505
+MYPRODDATA	+MYPRODDATA1479400505

Do not rename

Select this option if you do not want to rename mount points or ASM diskgroups during recovery. IBM Storage Copy Data Management will mount them with the same path/name as the source.

Add a custom prefix

Select this option and specify a custom prefix to be prepended to the source paths/names. The prefix value may contain leading or trailing slashes. In the case of ASM diskgroup names, the slashes are removed. For example:

<i>Table 11. Mount Point Rename - Add a custom prefix</i>		
Original Mount Point	Prefix	New Mount Point
/u02/myproddb/data	tmp	/tmpu02/myproddb/data
/u02/myproddb/data	/tmp/	/tmp/u02/myproddb/data
+MYPRODDATA	/tmp/	+TMPMYPRODDATA

Add a custom suffix

Select this option and specify a custom suffix to be appended the source paths/names. For example:

<i>Table 12. Mount Point Rename - Add a custom suffix</i>		
Original Mount Point	Suffix	New Mount Point
/u02/myproddb/data	tmp	/u02/myproddb/datatmp
/u02/myproddb/data	/tmp/	/u02/myproddb/data/tmp
+MYPRODDATA	/tmp/	+MYPRODDATATMP

Database Initialization Parameters

This option controls the initialization parameters used to start up the recovered database in Oracle Instant Database Restore and DevOps workflows.

Use same parameters as source

This is the default option. IBM Storage Copy Data Management uses the same initialization parameters as the source database, but with the following changes:

- Parameters that contain paths such as **control_files**, **db_recovery_file_dest**, or **log_archive_dest_*** are updated to reflect the new paths based on the renamed mount points of the recovered volumes.
- Parameters such as **audit_file_dest** and **diagnostic_dest** are updated to point to the appropriate location under the Oracle Base directory on the destination server if the path differs from the source server.
- The **db_name** and **db_unique_name** are updated to reflect the new name of the database if a new name is specified.
- Cluster-related parameters such as **instance_number**, **thread**, and **cluster_database** are set automatically by IBM Storage Copy Data Management depending on the appropriate values for the destination.

Use a template pfile

You can customize the initialization parameters by specifying a template file containing the initialization parameters that IBM Storage Copy Data Management should use.

The specified path must be to a plain text file that exists on the destination server and is readable by the IBM Storage Copy Data Management user. The file must be in Oracle pfile format, consisting of lines in the form `name = value`. Comments beginning with the `#` character are ignored.

IBM Storage Copy Data Management reads the template pfile and copies the entries to the new pfile that will be used to start up the recovered database. However, the following parameters in the template are ignored. Instead, IBM Storage Copy Data Management sets their values to reflect appropriate values from the source database or to reflect new paths based on the renamed mount points of the recovered volumes.

- `control_files`
- `db_block_size`
- `db_create_file_dest`
- `db_recovery_file_dest`
- `log_archive_dest`
- `spfile`
- `undo_tablespace`

Additionally, cluster-related parameters like **instance_number**, **thread**, and **cluster_database** are set automatically by IBM Storage Copy Data Management depending on the appropriate values for the destination.

Related tasks

[“Creating an Oracle Restore job definition” on page 157](#)

The procedures describe how to create an Oracle Restore job definition.

[“Creating an InterSystems Database Restore job definition” on page 151](#)

The procedures describe how to create an InterSystems Database Restore job definition. You can create InterSystems Caché and InterSystems IRIS application servers in IBM Storage Copy Data Management as an InterSystems Database.

Using state and status arguments in postscripts

The functionality described in this topic (**state** and **status** arguments) applies only to postscripts for Restore jobs. Backup jobs are not supported. Prescripts are not supported.

Certain positional arguments can be passed to a Restore job postscript for conditional logic. The arguments that can be passed are **state** and **status**. You can use the **status** argument, for example, if you want to perform an action only if a Restore job completes successfully. Your script would perform the action only if the value of **status** were SUCCESS.

When passing the arguments via the **Postscript** field in the IBM Storage Copy Data Management user interface, the arguments must be surrounded by underscores (`_STATE_` and `_STATUS_`). When doing so, the arguments are replaced in the script with the actual corresponding values.

Specifically, to pass the arguments **state** and **status**, enter the following in the **Parameters** field of the Postscript section of a job:

```
_STATE_ _STATUS_
```

So for the above postscript example, the following output would display if the job ran in an Instant Disk Restore state and completed successfully:

```
state IA
status SUCCESS
```

Note:

Passing `_STATE_` or `_STATUS_` to a script for an unsupported script type (such as prescripts) or an unsupported job type (such as Backup job scripts), simply passes “`_STATE_`” or “`_STATUS_`” as plain text.

An additional argument, `_VOLUMES_`, can also be added to a postscript to list the names of restored volumes.

Related concepts

[“Creating an IBM Storage Virtualize Restore job definition” on page 169](#)

[“Creating an NetAPP ONTAP Restore job definition” on page 179](#)

[“Creating a VMware Restore job definition” on page 191](#)

System jobs

The topics in the following section cover System job definitions as well as Maintenance job information.

Maintenance job

The Maintenance job removes resources and associated objects created by IBM Storage Copy Data Management when a job in a pending state is deleted. The cleanup procedure reclaims space on your storage devices, cleans up your IBM Storage Copy Data Management catalog, and removes related snapshots. By default, the Maintenance job runs once a day, but the job's associated schedule can be altered to run more or less frequently depending on your needs, or the job can be run manually. The job cannot be deleted.

The Maintenance job only performs cleanup operations once a job in a pending state is deleted. All logs associated with the deleted job are removed from IBM Storage Copy Data Management, so it is advised to download job logs before the Maintenance job's next run. The job can be stopped and resumed; all pending operations set to occur before the job was stopped will resume upon the next job run.

After deleting a pending Application, IBM, NetApp ONTAP, Pure Storage FlashArray, or VMware Backup job, all associated copy data, including recovery points, are deleted. The Maintenance job removes all VM Copies and Primary copies associated with deleted VMware Backup and Restore jobs. Similarly, after deleting a pending IBM, NetApp ONTAP, or Pure Storage FlashArray Backup or Restore job, all associated IBM, NetApp ONTAP, and Pure Storage FlashArray locations are removed by the Maintenance job. Once the Maintenance job completes, application, IBM, NetApp ONTAP, Pure Storage FlashArray or VMware data that was copied as part of the backup job cannot be recovered. Any data related to the deleted job will not be recoverable.

The Maintenance job also removes cataloged data associated with deleted Application, IBM, NetApp ONTAP, Pure Storage FlashArray and VMware Inventory jobs, and removes jobs and job sessions related to Script and Report jobs from the IBM Storage Copy Data Management interface.

Related tasks

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

Creating a report job definition

The procedures describe how to create a report job definition.

Before you begin

- If you have a specific set of data that you want to report on, create a report with customized parameters to include in the Report job. See [Create a Customized Report](#).
- At least one provider must be associated with an Inventory job. Before defining an Inventory job, add providers. See [Register a Provider](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).
- One or more schedules might also be associated with a job. Job sessions run based on the triggers defined in the schedule. See [Create a Schedule](#).

About this task

A Report job is a user-defined set of tasks and rules which run predefined or customized reports through a schedule that you define. The reports summarize information about cataloged providers and the data and other resources that reside on them. Reports generated during the job can be emailed in a variety of formats.

Procedure

1. Click the **Jobs** tab. Expand the **System** folder, then select **Reports**.
2. Click **New**, then select **Analyze**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. From the list of available reports, select one or more reports to include in the job definition. Expand reports to view associated customized reports. Select **Default** to run the predefined report parameters.
5. To create the job definition using default options, click **Create Job**. The job can be run manually from the Jobs tab.
6. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

7. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

8. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

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

What to do next

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

Related concepts

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Creating a script job definition

The procedures describe how to create a script job definition.

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 Storage Copy Data Management. See [Create a Schedule](#) and [Register a Provider](#).

About this task

A Script job defines a set of commands to run on the IBM Storage Copy Data Management appliance. Use the Script job feature to add functionality to IBM Storage Copy Data Management.

A Script job includes a user-defined script as well as other job parameters.

A script can consist of one or many commands, such as a shell script. Command entries may include any entry that is valid for a Linux CentOS command shell. The author of the script must be familiar with the command, options, and parameters.

As the script runs, IBM Storage Copy Data Management interprets a return code of 0 as success and subsequent values as follows:

- Return codes 1-125 indicate a failure, the meaning of which is specific to the command or shell script invoked
- Return code 126 indicates the script command was found but is not executable
- Return code 127 indicates that the script command was not found

- Return codes 128-255 indicate that the command abnormally terminated. The formula 128+N is used with N representing the signal the process terminated on
- Values outside the range of 0-255, in both negative and positive values, are wrapped around the same series of return codes and are matched to a corresponding code. When positive error codes display over 255, they read as follows: 256=0, 257=1, 258=2, etc. When negative return codes display, the corresponding codes are as follows: -1 = 255, -2 = 254, -3 = 253, etc.

Example of a simple script:

Following is a sample one line script that runs a single command:

```
echo hello
```

Once this script is saved and subsequently run as a job, the word “hello” is output on the virtual appliance.

Procedure

1. Click the **Jobs** tab. Expand the **System** folder, then select **Scripts**.
2. Click **New**, then select **Script**. The job editor opens.
3. Enter a name for your job definition and a meaningful description.
4. Enter the first command of your script in the command text box, and optionally enter a tag or remark in the corresponding description text box. Then click **Add**.
5. Continue to add commands. Use the reorder, edit, and delete options to assist you. Your script can consist of one or many commands.
6. To create the job definition using default options, click **Create Job**. The job can be run manually from the Jobs tab.
7. To edit options before creating the job definition, click **Advanced**. Set the job definition options.

Run scripts in order

Select **Run scripts in order** for the commands in the job to run sequentially. The job runs the first command and when it finishes it runs the next one. Clear this option for the commands to run concurrently.

Stop execution on failure

Select **Stop execution on failure** for the job to stop running as soon as one of the commands fails. Clear this option for the job to continue to run after one of the commands fails. This option is only applicable if **Run scripts in order** is selected.

8. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select the SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Email Address

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

9. Optionally, expand the **Schedule** section to select the job scheduling options. Select **Start job now** to create a job definition that starts the job immediately. Select **Schedule job to start at later time** to view the list of available schedules. Optionally select one or more schedules for the job. As each schedule is selected, the schedule's name and description displays.

Tip: To create and select a new schedule, click the **Configure** tab, then select **Schedules**. Create a schedule, return to the job editor, refresh the Available Schedules pane, and select the new schedule.

10. When you are satisfied that the job-specific information is correct, click **Create Job**. The job runs as defined by your schedule, or can be run manually from the **Jobs** tab.

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

What to do next

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

Related tasks

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

[“Creating a schedule” on page 87](#)

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

Related reference

[“Return code reference” on page 342](#)

Return Codes are issued when a script generated by a user-defined script is run. As the script runs, IBM Storage 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.

Editing a job definition

The procedures describe how to edit a job definition.

Before you begin

- Review the properties of your current job definitions. See [Jobs Overview](#).

About this task

Revise a job definition to change the provider that the job is run against, SLA Policy associated with a job definition, options, notification and schedule properties.

Procedure

1. Click the **Jobs** tab.
2. Select the job definition to edit by clicking in the row containing the job definition name.
3. Click **Edit**. The Job Definition Editor opens.
4. Click through the wizard, making revisions as needed.
5. Click **Finish** when you are satisfied that the job-specific information is correct.

What to do next

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

Related tasks

[“Deleting a job definition” on page 212](#)

The procedures describe how to delete a job definition.

Deleting a job definition

The procedures describe how to delete a job definition.

About this task

Delete a job definition when it becomes obsolete. This keeps your operations current.

Procedure

1. Click the **Jobs** tab.
2. Select the job definition to delete by clicking in the row containing the job definition name.
3. Click **Delete**. A confirmation dialog box opens.
4. Confirm deletion. The job definition is deleted.

Related concepts

[“Maintenance job” on page 207](#)

Related tasks

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

Chapter 16. Searching

The topics in the following section cover searching for objects, downloading search results, and browsing the Inventory.

Searching overview

Use IBM Storage 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 IBM Inventory may include FlashCopies, Hosts, IOGroups, MDisks, mirrors, Node Canisters, PortIPs, and volumes.

The NetApp ONTAPInventory may include aggregates, CIFS shares, files, LUNs, networks, NFS exports, nodes, policies, protocols, qtrees, quotas, SnapMirrors, Snapshots, SnapVaults, SVMs, vFilers, and volumes.

The Recovery Inventory may contain datacenters, data stores, ESX hosts, LUNs, folders, recovery points, vApps, vDisks, vSnapshots, and vSpheres.

The VMware Inventory may include data stores, ESX hosts, LUNs, virtual disks, virtual machines, VMware hosts, and virtual snapshots.

The Pure Storage FlashArray Inventory may include volumes, snapshots, LUNS, hosts, and host groups.

You can match a character pattern and apply other filters such as category, object type, and location through the Search feature. The results are presented in the user interface and are also exportable. Furthermore, click an object that appears in the Search results to open a tab with additional details about that object. You can also review previous versions of your files, along with their Snapshot, SnapVault and SnapMirror replication status.

Alternatively, use the Inventory Browser to browse through the list of providers. Drill into the Inventory Browser to logically view the details of the objects underlying a storage system, virtual host, or application.

Use the Time Machine control to view the Inventory as it appeared on a past date.

Tip: IBM Storage Copy Data Management lets you quickly and easily locate every version of a file across your entire Enterprise. IBM Storage 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 concepts

[“Browsing inventory” on page 222](#)

Use IBM Storage 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.

Related tasks

[“Viewing object details” on page 218](#)

From the search results, you can view the attributes of a searched object, including its location, type, and the dates associated with its creation and modification. If more than one version of an object exists in the Inventory, the search results displays the latest version and attributes of the object.

[“Viewing NetApp ONTAP file details” on page 219](#)

From the search results, you can view the attributes of a searched file, including its location, size, and the date the file was added to the Inventory. You can also quickly view a file's creation time, last modified time, and the last accessed time from the search results pane.

[“Finding and restoring a file” on page 220](#)

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 Storage Copy

Data Management returns the location of every instance and version of that file across your NetApp ONTAP storage infrastructure.

[“Downloading search results” on page 221](#)

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.

Related reference

[“Searching for objects” on page 214](#)

Use the search feature to find objects on providers that are cataloged in IBM Storage Copy Data Management. Examples of objects are volumes, files, snapshots, qtrees, and virtual machines. You can tailor your search by applying filters.

Searching for objects

Use the search feature to find objects on providers that are cataloged in IBM Storage 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 Storage Copy Data Management searches the entire Inventory and returns all objects with a name that matches or contains the search entry.

Advanced search is similar to basic search, with an additional function. You can search and filter by object name, category, object type, and location. When searching for NetApp ONTAP files, you can also filter by last modified time, creation time, last accessed time, and file size.

Once you initiate the search and it completes, IBM Storage 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 [“Registering a provider” on page 25](#) and Plan Overview.

To search for objects:

1. Click the **Search** tab.
2. Open a new Search pane. If this is the first search in your IBM Storage 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, 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.

What to do next

- You can download the search results as a CSV file format. See [“Downloading search results” on page 221](#).
- 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 [“Viewing object details” on page 218](#).

Related concepts

[“Browsing inventory” on page 222](#)

Use IBM Storage 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.

Related tasks

[“Downloading search results” on page 221](#)

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.

[“Viewing object details” on page 218](#)

From the search results, you can view the attributes of a searched object, including its location, type, and the dates associated with its creation and modification. If more than one version of an object exists in the Inventory, the search results displays the latest version and attributes of the object.

[“Viewing NetApp ONTAP file details” on page 219](#)

From the search results, you can view the attributes of a searched file, including its location, size, and the date the file was added to the Inventory. You can also quickly view a file's creation time, last modified time, and the last accessed time from the search results pane.

Related reference

[“Search and filter guidelines” on page 339](#)

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.

[“Select, sort, and reorder columns” on page 341](#)

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.

Viewing object details

From the search results, you can view the attributes of a searched object, including its location, type, and the dates associated with its creation and modification. If more than one version of an object exists in the Inventory, the search results displays the latest version and attributes of the object.

Before you begin

- You can only browse objects on providers that are registered and cataloged.
- Search for an object. Note that you can filter search results to only return specific objects by using an advanced search.

Restriction: Note that this topic does not apply to searches for low-level NetApp ONTAP files cataloged with a NetApp ONTAP File Inventory job. Results returned from a NetApp ONTAP file search differ from other object search results.

Procedure

To view object details, complete the following steps:

1. Click an object name in the search results pane to view more information about an object. The properties of the object display in a new tab.
2. If more than one version of the object exists in the Inventory, select a version of the object to view from the Versions tab.

Note: The object that displays in the search results pane is the latest version. In some cases, the name of a returned object on the search results pane may be the same as another object, however the resources where the objects reside is different. Review the file properties of the objects by selecting their names on the search results pane to view the differences between the returned entries.

3. Review the properties of the object. The specific properties vary by type of object.

Related concepts

[“Searching overview” on page 213](#)

Use IBM Storage 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.

Related tasks

[“Viewing NetApp ONTAP file details” on page 219](#)

From the search results, you can view the attributes of a searched file, including its location, size, and the date the file was added to the Inventory. You can also quickly view a file's creation time, last modified time, and the last accessed time from the search results pane.

Related reference

[“Searching for objects” on page 214](#)

Use the search feature to find objects on providers that are cataloged in IBM Storage Copy Data Management. Examples of objects are volumes, files, snapshots, qtrees, and virtual machines. You can tailor your search by applying filters.

Viewing NetApp ONTAP file details

From the search results, you can view the attributes of a searched file, including its location, size, and the date the file was added to the Inventory. You can also quickly view a file's creation time, last modified time, and the last accessed time from the search results pane.

Before you begin

- You can only view all available file details once your NetApp ONTAP providers are cataloged through both a NetApp ONTAP Storage Inventory job and a NetApp ONTAP File Inventory job. By enabling the Catalog all available snapshots option in the NetApp ONTAP File Inventory job, you can view multiple versions of a file. See [“Creating a NetApp ONTAP Storage Inventory job definition” on page 103](#), and [“Creating a NetApp ONTAP File Inventory job definition” on page 104](#).
- Search for a file by selecting **NetApp > File** from the Search For filter on the Advanced Search dialog. See [“Searching for objects” on page 214](#).

About this task

Note: this topic only applies to searches for low-level NetApp ONTAP files cataloged with a NetApp ONTAP File Inventory job.

On the searched file's properties pane, you can review previous versions of your files, along with their Snapshot, SnapVault, and SnapMirror replication status on the file's properties pane.

If more than one version of a file exists in the Inventory, the search results pane displays the latest version and attributes of the file. View the file's properties to view previous versions and associated attributes of the file.

Note that you can filter search results to only return files through an advanced search.

Procedure

To view file details, complete the following steps:

1. Click a file name in the search results pane to view more information about a file. The properties of the file display in a new tab.
2. If more than one version of the file exists in the Inventory, select a version of the file to view from the Versions tab.

Note: The file displayed in the search results pane is the latest version. In some cases, the name of a returned object on the search results pane may be the same as another object, however the resources where the objects reside is different. Review the file properties of the objects by selecting their names on the search results pane to view the differences between the returned entries.

3. Click an instance in the Snapshots, SnapMirrors, or SnapVaults tab to view details about the associated replication instance, such as the creation date, mirror or vault location, and the source.

What to do next

Restore a file to a previous version. See [“Finding and restoring a file” on page 220](#).

Related concepts

[“Searching overview” on page 213](#)

Use IBM Storage 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.

Related tasks

[“Viewing object details” on page 218](#)

From the search results, you can view the attributes of a searched object, including its location, type, and the dates associated with its creation and modification. If more than one version of an object exists in the Inventory, the search results displays the latest version and attributes of the object.

[“Finding and restoring a file” on page 220](#)

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 Storage Copy Data Management returns the location of every instance and version of that file across your NetApp ONTAP storage infrastructure.

Related reference

[“Searching for objects” on page 214](#)

Use the search feature to find objects on providers that are cataloged in IBM Storage Copy Data Management. Examples of objects are volumes, files, snapshots, qtrees, and virtual machines. You can tailor your search by applying filters.

Finding and restoring 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 Storage 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 [“Creating a NetApp ONTAP File Inventory job definition” on page 104](#).
- To enable NFS links from the **File Links** tab on an object's property pane, access IBM Storage Copy Data Management through a browser in a Linux environment. Additionally, an automounter must be running using the default local path of /net.
- By default, CIFS users cannot see the .snapshot directory. Set the **cifs.show_snapshot** option to **On** to view the .snapshot directory.
- Activate snapshot functionality on the host and enable the necessary access privileges for the user accessing the host. See [Use a Script Job to Run Snap Creator](#).

Procedure

To find and restore a file, complete the following steps:

To restore a file on a Windows machine, search for the file in IBM Storage 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:
 - a. Select a previous version of the file and click **Open** to open the previous version without altering the original version.

- b. Select a previous version of the file and click **Copy** to copy the previous version to an alternate location.
- c. Select a previous version of the file and click **Restore** to replace or roll back the file to the selected previous version.

Related concepts

[“Searching overview” on page 213](#)

Use IBM Storage 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.

Related tasks

[“Viewing NetApp ONTAP file details” on page 219](#)

From the search results, you can view the attributes of a searched file, including its location, size, and the date the file was added to the Inventory. You can also quickly view a file's creation time, last modified time, and the last accessed time from the search results pane.

Related reference

[“Searching for objects” on page 214](#)

Use the search feature to find objects on providers that are cataloged in IBM Storage Copy Data Management. Examples of objects are volumes, files, snapshots, qtrees, and virtual machines. You can tailor your search by applying filters.

Downloading 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.

Before you begin

- Use the search feature to find objects on providers that are cataloged in IBM Storage Copy Data Management. See [“Searching for objects” on page 214](#).

Note: Hyperlinks might not be active in the CSV file. Only search results that are visible on the search results pane are exported.

Procedure

To save search results as a CSV file, complete the following steps:

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 tasks

[“Viewing object details” on page 218](#)

From the search results, you can view the attributes of a searched object, including its location, type, and the dates associated with its creation and modification. If more than one version of an object exists in the Inventory, the search results displays the latest version and attributes of the object.

[“Viewing NetApp ONTAP file details” on page 219](#)

From the search results, you can view the attributes of a searched file, including its location, size, and the date the file was added to the Inventory. You can also quickly view a file's creation time, last modified time, and the last accessed time from the search results pane.

Related reference

[“Searching for objects” on page 214](#)

Use the search feature to find objects on providers that are cataloged in IBM Storage Copy Data Management. Examples of objects are volumes, files, snapshots, qtrees, and virtual machines. You can tailor your search by applying filters.

Browsing inventory

Use IBM Storage 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 IBM Inventory may include FlashCopies, Hosts, IOGroups, MDisks, mirrors, Node Canisters, PortIPs, and volumes.

The NetApp ONTAPInventory may include aggregates, CIFS shares, files, LUNs, networks, NFS exports, nodes, policies, protocols, qtrees, quotas, SnapMirrors, Snapshots, SnapVaults, SVMs, vFilers, and volumes.

The Recovery Inventory may contain datacenters, data stores, ESX hosts, LUNs, folders, recovery points, vApps, vDisks, vSnapshots, and vSpheres.

The VMware Inventory may include data stores, ESX hosts, LUNs, virtual disks, virtual machines, VMware hosts, and virtual snapshots.

The Pure Storage FlashArray Inventory may include volumes, snapshots, LUNS, hosts, and host groups.

Browsing cataloged providers

Before you begin

You can only browse objects on providers that are registered and cataloged.

Procedure

1. Click the **Search** tab. On the Views pane, select **Catalog**. The Inventory Browser opens.
2. Drill down through providers in the Inventory Browser. A tab that displays the names of the underlying objects opens.
3. In the tab, click an object name. The properties of the object display in a new tab. The specific properties vary by type of object.

Related concepts

[“Searching overview” on page 213](#)

Use IBM Storage 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.

Related tasks

[“Viewing object details” on page 218](#)

From the search results, you can view the attributes of a searched object, including its location, type, and the dates associated with its creation and modification. If more than one version of an object exists in the Inventory, the search results displays the latest version and attributes of the object.

[“Viewing NetApp ONTAP file details” on page 219](#)

From the search results, you can view the attributes of a searched file, including its location, size, and the date the file was added to the Inventory. You can also quickly view a file's creation time, last modified time, and the last accessed time from the search results pane.

Related reference

“Searching for objects” on page 214

Use the search feature to find objects on providers that are cataloged in IBM Storage Copy Data Management. Examples of objects are volumes, files, snapshots, qtrees, and virtual machines. You can tailor your search by applying filters.

Browsing through the Inventory on a previous date

Before you begin

You can only browse objects on providers that are registered and cataloged.

Procedure

To browse through the Inventory on a previous date, complete the following steps:

1. Click the **Search** tab. On the Views pane, select **Catalog**. The Inventory Browser opens.
2. Click the **All Resources** tab. The Time Machine control appears.
3. Use the calendar tool to set the Time Machine to a previous date. The Inventory Browser reflects the state of the Inventory on that date.
4. Drill down through providers in the Inventory Browser. A tab that displays the names of the underlying objects opens.

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

Chapter 17. Report

The topics in the following section cover running and customizing reports, as well as individual report details.

Report Overview

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

The information in these reports are presented in a chart-based **Quick View** section, or tabular **Summary View** and **Detail View** sections.

Reports include interactive elements, such as searching for individual values within a report, vertical scrolling, and column sorting. Information groups, such as the Primary Source Volume groups in the NetApp ONTAP Protection Usage report, can also be sorted by clicking the group name.

You can add a Report job to summarize information about cataloged providers and the data and other resources that reside on them, then schedule the Report job to run as defined by the parameters of the schedule.

To further analyze the data or print a hard copy, use the export functionality to save the data from the generated report to an Adobe PDF, Microsoft Word file, Microsoft Excel file, or HTML file.

Related concepts

[“Application reports” on page 230](#)

The Application Reports help you review your application server configuration. Use the Application Reports to view your application server's database, log disks, and eligibility for protection. Reports are based on the data collected by the most recently run job.

[“System management reports” on page 236](#)

System Management Reports offer an in-depth view of the status of your IBM Storage Copy Data Management configuration, including cataloged storage system information, jobs, and their status.

[“File analytics reports” on page 243](#)

The File Analytics Reports help you review your storage needs and examine your storage capacity. There are a number of variables that you can use to view the files on your system. Using the File Analytics Reports you can drill down to information about the size of files, the age of files, and type of files on your storage systems. Reports are based on the data collected by the most recently run job.

[“Protection compliance reports” on page 250](#)

The Protection Compliance reports help ensure your data is protected through user-defined recovery point objective parameters.

[“Storage protection reports” on page 269](#)

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.

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

[“Editing a customized report” on page 228](#)

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.

[“Downloading a Report” on page 229](#)

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.

[“Creating a report job definition” on page 208](#)

The procedures describe how to create a report job definition.

Running a report

Perform the following steps to run any report from the Report tab. You can run reports with predefined default parameters or run customized reports driven by custom parameters.

Before you begin

- At least one provider must be associated with a job definition. See [“Registering a provider” on page 25](#).
- Run a job immediately; sometimes referred to as on demand. See [“Start, pause, and hold a job session” on page 91](#).
- You can also create a schedule that can be applied to a job definition so that the job is run as defined by the parameters of the schedule. See [“Creating a schedule” on page 87](#).

Procedure

To run a report, complete the following steps:

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. To run a report for default and custom parameters, complete the following steps:
 - **For default parameters:**
Click **Run**.The default report data is returned in the Report pane.
 - **For custom parameters:**
Select report parameter values in the **Parameters** pane and then click **Run**. Parameters are unique to each report. The parameters that you select drive the report output. The customized report data is returned in the Report pane.

What to do next

- Create a Report job definition to schedule the report to run automatically. See [“Creating a report job definition” on page 208](#).
- Save a report with customized parameters. See [“Creating a customized report” on page 227](#)

Related concepts

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Creating a customized report” on page 227](#)

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.

[“Creating a report job definition” on page 208](#)

The procedures describe how to create a report job definition.

[“Downloading a Report” on page 229](#)

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.

Creating a customized report

Perform the following steps to create a report with customized parameters. Select a predefined report, set custom parameters, and save the report with a customized name to run on demand or create a schedule to run the report as defined by the parameters of the schedule. Customized reports display nested under their predefined source report on the Report Browser pane.

Before you begin

- At least one provider must be associated with a job definition. See [“Registering a provider” on page 25](#).

Procedure

To create a customized report, complete the following steps:

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.

What to do next

Run the customized report from the **Report** tab. See [“Running a report” on page 226](#).

Related concepts

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Editing a customized report” on page 228](#)

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.

[“Running a report” on page 226](#)

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.

[“Downloading a Report” on page 229](#)

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.

Editing a customized report

Perform the following steps to edit a report with customized parameters. Edit the custom parameters, and save the report with a customized name to run on demand or create a schedule to run the report as defined by the parameters of the schedule. Customized reports display nested under their predefined source report on the Report Browser pane.

Before you begin

- At least one provider must be associated with a job definition. See [“Registering a provider” on page 25](#).
- Run a job immediately; sometimes referred to as on demand. See [“Start, pause, and hold a job session” on page 91](#).
- Create a customized report. See [“Creating a customized report” on page 227](#).

Procedure

To edit a customized report, complete the following steps:

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.

What to do next

Run the customized report from the **Report** tab. See [“Running a report” on page 226](#).

Related concepts

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Creating a customized report” on page 227](#)

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.

[“Running a report” on page 226](#)

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.

[“Downloading a Report” on page 229](#)

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.

Downloading 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.

Before you begin

Run a report from the Report Browser pane. See [“Running a report” on page 226](#).

About this task

Note: 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.

Some report images appear truncated when viewed in Microsoft Word. For best results, view downloaded reports through the Web Layout.

Procedure

To download a report, complete the following steps:

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.
4. To download a generated report, complete the following steps:
 - a) Click the **Report** tab. On the Views pane, select **Jobs History**.
 - b) On the All Reports pane, select a generated report to download.
 - c) Click the **HTML** , **PDF** , **Excel** , or **Word** icon associated with the report. You can view the report now or save it to a file.

Reports are saved on the IBM Storage Copy Data Management appliance in the following directories:

- Report Download location: /data/reports/output
- Report Name location: /data/reports/repository/reports/output
- Report Temporary location: /data/reports/repository/reports/doc

What to do next

Review your downloaded report for further analysis of your data.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

Deleting a generated report

Delete generated reports from the list of generated reports on the Jobs History pane.

Before you begin

Note: 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.

Procedure

To delete a generated report, complete the following steps:

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.

What to do next

Run the customized report from the **Report** tab. See [“Running a report” on page 226](#).

Related tasks

[“Running a report” on page 226](#)

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.

[“Downloading a Report” on page 229](#)

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.

Application reports

The Application Reports help you review your application server configuration. Use the Application Reports to view your application server's database, log disks, and eligibility for protection. Reports are based on the data collected by the most recently run job.

Choose the Application reports that fit your needs:

- **Application Configuration Report** - Review the configuration of your application servers including the disks you the database and logs reside.
- **Application RPO Compliance Report** - The Application RPO Compliance report displays your application servers in relation to your recovery point objective parameters.

Related concepts

[“Application configuration report” on page 231](#)

Review the configuration of your application servers including the disks where the database and logs reside.

[“Application RPO Compliance report” on page 233](#)

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.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database

condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

[“Downloading a Report” on page 229](#)

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.

Application configuration report

Review the configuration of your application servers including the disks where the database and logs reside.

Before you begin

- Create and run a Database Inventory job. You can select one or more application provider in a single job definition for cataloging. See [“Creating a Database Inventory job definition” on page 97](#).
- For Microsoft-based application servers, ensure the vCenter containing the virtual machines are registered in IBM Storage Copy Data Management. Ensure that the virtual machine credentials for the virtual machines hosting the applications are configured. See [“Adding a credential” on page 70](#).

Parameters

Use the following parameters to customize your report:

- Application Type: File System, InterSystems Caché and InterSystems IRIS (collectively referred to as InterSystems Database), Microsoft SQL, Oracle, SAP HANA. Multiple selections are supported.

The default report parameter, All, reports on all available application types in your configuration.

Detail view - Microsoft SQL

The following fields and corresponding data display in the Microsoft Active Directory, Microsoft Exchange, Microsoft SQL, and Microsoft SharePoint sections of the Application Configuration report:

Instance

The name and location of the SQL instance.

Database

The name of the database associated with the SQL instance and application server.

Data Disk(s)

The name of the associated database disk mount points, along with the associated volume and storage array in parentheses.

Log Disk(s)

The name of the associated log disk mount points, along with the associated volume and storage array in parentheses.

Eligible for Protection

The status of the database's eligibility for protection.

Detail view - Oracle

The following fields and corresponding data display in the Oracle section of the Application Configuration report:

Oracle Home

The name and location of the Oracle home.

Database

The name of the database associated with the Oracle home and application server.

Data Disk(s)

The name of the associated database disk mount points, along with the associated volume and storage array in parentheses.

Log Disk(s)

The name of the associated log disk mount points, along with the associated volume and storage array in parentheses.

Eligible for Protection

The status of the database's eligibility for protection.

Detail view - InterSystems Database

The following fields and corresponding data display in the InterSystems Database section of the Application Configuration report:

Instance

The name and location of the InterSystems Database instance.

Database

The name of the database associated with the InterSystems Database instance and application server.

Data Disk(s)

The name of the associated database disk mount points, along with the associated volume and storage array in parentheses.

Log Disk(s)

The name of the associated log disk mount points, along with the associated volume and storage array in parentheses.

Eligible for Protection

The status of the database's eligibility for protection.

Detail view - SAP HANA

The following fields and corresponding data display in the SAP HANA section of the Application Configuration report:

Instance

The name and location of the SAP HANA instance.

Database

The name of the database associated with the SAP HANA instance and application server.

Data Disk(s)

The name of the associated database disk mount points, along with the associated volume and storage array in parentheses.

Log Disk(s)

The name of the associated log disk mount points, along with the associated volume and storage array in parentheses.

Eligible for Protection

The status of the database's eligibility for protection.

Detail view - File System

The following fields and corresponding data display in the File System section of the Application Configuration report:

File System

The name of the database associated with the file system.

Mount Points

The name of the associated database disk mount points, along with the associated volume and storage array in parentheses.

File System Type

The type of registered file system.

Eligible for Protection

Eligible for Protection

The status of the database's eligibility for protection.

Related concepts

[“Application reports” on page 230](#)

The Application Reports help you review your application server configuration. Use the Application Reports to view your application server's database, log disks, and eligibility for protection. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

Application RPO Compliance report

The Application RPO Compliance report displays your application servers in relation to your recovery point objective parameters. Determine which of your application servers are not in compliance with your RPO parameters, and discover the reasons for their non-compliance.

Before you begin

- Create and run a Database Inventory job. You can select one or more application provider in a single job definition for cataloging. See [“Creating a Database Inventory job definition” on page 97](#).
- Create and run a Database Backup job.

Parameters

Use the following parameters to customize your report:

- Application Type

InterSystems Caché and InterSystems IRIS (collectively referred to as InterSystems Database), Microsoft SQL, Oracle, SAP HANA. Multiple selections are supported.

- Application Server

Multiple selections are supported.

- Protection Type

Set the protection types to return in the report. Values include Primary and Replication. Multiple selects are supported.

- Storage Vendor

Set the storage vendor types to display in the report. Multiple selections are supported.

- Display Databases That Are

Set the compliance status of your application server to return in the report. Values include Compliant and Not Compliant. Multiple selections are supported. By default, this parameter is set to Not Compliant.

- RPO Older Than

Set the age of the recovery point objective in days.

The default report parameters report on all non-compliant application servers based on an RPO older than one day.

Quick View

The Quick View section displays a bar graph of compliant and non-compliant application servers based on your RPO parameters.

Not Compliant for Primary Protection

The following fields and corresponding data display in the Not Compliant for Primary Protection section of the Application RPO Compliance report:

Database

The names and properties of the resource.

Application Server (Type)

The application server's location and type.

Application Instance

The name of the application instance, or Oracle home if reporting on Oracle databases.

Job Name

The Backup job associated with the application server.

Storage Vendor

The storage vendor associated with the application server.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Reason

The reason the application server does not meet your RPO compliance parameters. Examples include no successful runs of a backup job, or backing up to an unsupported disk.

Compliant for Primary Protection

The following fields and corresponding data display in the Compliant for Primary Protection section of the Application RPO Compliance report:

Database

The names and properties of the resource.

Application Server (Type)

The application server's location and type.

Application Instance

The name of the application instance, or Oracle home if reporting on Oracle databases.

Job Name

The Backup job associated with the application server.

Storage Vendor

The storage vendor associated with the application server.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your application server will be non-compliant.

Not Compliant for Replication

The following fields and corresponding data display in the Not Compliant for Replication section of the Application RPO Compliance report:

Database

The names and properties of the resource.

Source

The source of the secondary protection.

Destination

The destination of the secondary protection.

Job Name

The Backup job associated with the application server.

Storage Vendor

The storage vendor associated with the application server.

Protection Time

The most recent instance of a successful run of the Backup job.

Reason

The reason the application server does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Replication

The following fields and corresponding data display in the Compliant for Replication section of the Application RPO Compliance report:

Database

The names and properties of the resource.

Source

The source of the secondary protection.

Destination

The destination of the secondary protection.

Job Name

The Backup job associated with the application server.

Storage Vendor

The storage vendor associated with the application server.

Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your application server will be non-compliant.

Related concepts

[“Application reports” on page 230](#)

The Application Reports help you review your application server configuration. Use the Application Reports to view your application server's database, log disks, and eligibility for protection. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

System management reports

System Management Reports offer an in-depth view of the status of your IBM Storage Copy Data Management configuration, including cataloged storage system information, jobs, and their status.

Use the System Management Reports to answer questions such as:

- What is the operating system, memory, and processor information of my storage system?
- How many jobs are associated with an Inventory?
- What is the average runtime of a job, and when was it last successfully run?

Choose the System Management report that fits your needs:

- **Catalog Summary Report** - Displays a summary of all the high level objects cataloged.
- **Configuration Report** - Review IBM Storage Copy Data Management node configuration and associated jobs.
- **Job Report** - Review jobs and their status, including the last time they ran successfully.
- **Job Sessions Report** - Review sessions for jobs and their status.

Related concepts

[“Catalog summary report” on page 237](#)

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.

[“Application configuration report” on page 231](#)

Review the configuration of your application servers including the disks where the database and logs reside.

[“Job report” on page 240](#)

Review the available jobs in your IBM Storage Copy Data Management configuration. Run the Job report to view jobs by type, their average runtime, and their successful run percentage.

[“Job sessions report” on page 242](#)

Review sessions for jobs and their status in your IBM Storage Copy Data Management environment. Run the Job Sessions report to view job sessions by type, their status, start time, finish time, and duration.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

[“Downloading a Report” on page 229](#)

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.

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, IBM Storage Virtualize, NetApp ONTAP, NetApp ONTAP Node Summary, Pure Storage FlashArray, and VMware. Multiple selections are supported.

The default report parameter, **All**, reports on all available Catalog information.

NetApp ONTAP Node Summary

The NetApp ONTAP Node Summary section of the report displays an overview of your NetApp ONTAP nodes, including the number of aggregates, volumes, snapshots, and the total number of files on your nodes.

NetApp ONTAP Objects, VMware Objects, Application Objects, IBM Storage Virtualize Objects, and Pure Storage FlashArray Objects

The following fields and corresponding data display in the NetApp ONTAP Objects, VMware Objects, Application Objects, and IBM Objects, and Pure Storage FlashArray Objects sections of the Catalog Summary report:

Object

The name of the cataloged high level object.

Total Count

The number of records found on the cataloged object.

Catalog Size

The total amount of data on the cataloged object.

Count from Last Update

The latest number of records cataloged during the last update.

Last Updated

The date and time of the most recent Catalog update.

Related concepts

[“System management reports” on page 236](#)

System Management Reports offer an in-depth view of the status of your IBM Storage Copy Data Management configuration, including cataloged storage system information, jobs, and their status.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

Configuration report

Review the IBM Storage Copy Data Management providers added to IBM Storage 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, Application Nodes, IBM Storage 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 Catalog information.

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 Storage Virtualize Nodes

The following fields and corresponding data display in the IBM Storage Virtualize Nodes section of the Configuration report:

Node (Site)

The name of the IBM Storage Virtualize node and the associated site.

Model

The Storage Virtualize storage model name.

Operating System Version

The installed operating system version on the IBM Storage Virtualize node.

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.

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 concepts

[“System management reports” on page 236](#)

System Management Reports offer an in-depth view of the status of your IBM Storage Copy Data Management configuration, including cataloged storage system information, jobs, and their status.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

Job report

Review the available jobs in your IBM Storage Copy Data Management configuration. Run the Job report to view jobs by type, their average runtime, and their successful run percentage.

Use the Job report to answer questions such as:

- What types of jobs are available?
- What is the average runtime of a job, and when was it last successfully run?
- How many times has a specific job run successfully or failed?

Parameters

Use the following parameters to edit your report:

- Job Type
 - Multiple selections are supported.
- Days Since Successful Run
- Show Job Details
 - Enable to display the Job ID field in the Detail View section.

The default report parameters report on all available jobs.

Quick View

The Quick View section displays a pie chart of the number of times a job type successfully completed, failed, or was marked with any of the following statuses: unknown, waiting, running, stopped, partial, skipped, aborted, or stopping. Use the Job Type parameter to display File, NetApp ONTAP, Report, Script, VMware, or all job types.

Note: The Quick View section is only modified through the Job Type parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the Jobs report:

Job Type

The type of job. For example, a catalog or report job.

Jobs

The number of jobs associated with the job type.

Runs

The number of times a job of this type ran.

Completed

The number of times a job of this type successfully completed.

Failed

The number of times a job of this type failed.

Other

The number of times a job of this type was marked with any of the following statuses: unknown, waiting, running, stopped, partial, skipped, aborted, or stopping.

Detail View

The following fields and corresponding data display in the Detail View section of the Jobs report:

Job

The job name.

Job ID

The job ID of the displayed job. This field is controlled by the Show Job Details parameter.

Type

The type of job. For example, a catalog or report job.

Runs

The number of times the job ran.

Completed

The number of times the job successfully completed.

Failed

The number of times the job failed.

Other

The number of times the job was marked with any of the following statuses: unknown, waiting, running, stopped, partial, skipped, aborted, or stopping.

Last Successful Run

The date and time the job last ran successfully.

Average Runtime (Days hh:mm:ss)

The average time it takes to run the job, listed in days, hours, minutes, and seconds.

Success %

The percentage of times the job successfully completed.

Related concepts

[“System management reports” on page 236](#)

System Management Reports offer an in-depth view of the status of your IBM Storage Copy Data Management configuration, including cataloged storage system information, jobs, and their status.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

Job sessions report

Review sessions for jobs and their status in your IBM Storage Copy Data Management environment. Run the Job Sessions report to view job sessions by type, their status, start time, finish time, and duration.

Use the Job report to answer questions such as:

- What job sessions have run in the environment?
- What are the types of job sessions that have run in the environment?
- What is the status of those job sessions?
- What was the start and end time for each job session?
- What was the duration from start to finish for each job session?

Parameters

Use the following parameters to edit your report:

- Job Type
Multiple selections are supported.
- Job
Multiple selections are supported.
- Status
Multiple selections are supported.
- Protection Type
Multiple selections are supported.
- Job Sessions for Past Number of Days
Enter a number for the number of days.

The default report parameters report on all available jobs and 1 for Job Sessions for Past Number of Days.

Detail view

The following fields and corresponding data display in the Summary View section of the Jobs report:

Job

The job name.

Type

The type of job. For example, a catalog or report job.

Status

The status of the job session.

Start Time

The time that the job session started in UTC time.

Finish Time

The time that the job session finished in UTC time.

Duration (Days hh:mm:ss)

The total time that it takes for the job session to run, listed in days, hours, minutes, and seconds.

Related concepts

[“System management reports” on page 236](#)

System Management Reports offer an in-depth view of the status of your IBM Storage Copy Data Management configuration, including cataloged storage system information, jobs, and their status.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

File analytics reports

The File Analytics Reports help you review your storage needs and examine your storage capacity. There are a number of variables that you can use to view the files on your system. Using the File Analytics Reports you can drill down to information about the size of files, the age of files, and type of files on your storage systems. Reports are based on the data collected by the most recently run job.

Use the File Analytics Reports to answer questions such as:

- What is the average size of files stored?
- What files have not been accessed within a given time period?
- What is the amount of space usage by file type?

Note: IBM Storage 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.

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.

Related concepts

[“File usage by owner report” on page 244](#)

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.

[“Files by age report” on page 246](#)

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.

[“Files by category report” on page 247](#)

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.

[“Files by size report” on page 249](#)

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.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

[“Downloading a Report” on page 229](#)

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.

File usage by owner report

the File Usage by Owner report to view the owners consuming the largest amount of space on your storage systems as well as the number of files associated with an owner.

Use the File Usage by Owner report to answer questions such as:

- Which owners are consuming the largest amount of space on a selected storage system?
- How many files are associated with a specific owner on a selected storage system?

BEFORE YOU BEGIN:

- Create and run a NetApp ONTAP File Inventory job with the Traversal Mode set to Filewalk in conjunction with the IBM Storage Copy Data Management Filewalker tool. See [“Creating a NetApp ONTAP File Inventory job definition” on page 104](#).

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 concepts

[“File analytics reports” on page 243](#)

The File Analytics Reports help you review your storage needs and examine your storage capacity. There are a number of variables that you can use to view the files on your system. Using the File Analytics Reports you can drill down to information about the size of files, the age of files, and type of files on your storage systems. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

Files by age report

Review the age of files on your node based on the creation date and the last time modified. Run the Files By Age report to view the age of files on your node, from older than ten year to less than one year.

Age categories include:

- Older than ten years
- Five to ten years
- Four to five years
- Three to four years
- Two to three years
- One to two years
- 180 days to one year
- Less than 180 days

Use the Files by Age report to answer questions such as:

- How much storage space and the percentage of the total space is used by files created more than ten years ago?
- How much storage space is used within a given time period?

BEFORE YOU BEGIN:

- To view accessed time information, create and run a NetApp ONTAP File Inventory job with the Traversal Mode set to Filewalk in conjunction with the IBM Storage Copy Data Management Filewalker tool. See [“Creating a NetApp ONTAP File Inventory job definition” on page 104](#).

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 concepts

[“File analytics reports” on page 243](#)

The File Analytics Reports help you review your storage needs and examine your storage capacity. There are a number of variables that you can use to view the files on your system. Using the File Analytics Reports you can drill down to information about the size of files, the age of files, and type of files on your storage systems. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

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 concepts

[“File analytics reports” on page 243](#)

The File Analytics Reports help you review your storage needs and examine your storage capacity. There are a number of variables that you can use to view the files on your system. Using the File Analytics Reports you can drill down to information about the size of files, the age of files, and type of files on your storage systems. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

Files by size report

Identify the largest space consumers on your storage systems to help manage storage utilization. Run the Files By Size report to view the largest files, the path and volume where they are located, and their size. View the amount of space the largest file on your storage system utilizes as a percentage of storage system capacity.

Use the Files by Size report to answer questions such as:

- What are the largest files on the system?
- What are the sizes of those files?
- What storage system do they reside on and its associated volume?

BEFORE YOU BEGIN:

- To view file ownership information, create and run a NetApp ONTAP File Inventory job with the Traversal Mode set to Filewalk in conjunction with the IBM Storage Copy Data Management Filewalker tool. See [“Creating a NetApp ONTAP File Inventory job definition” on page 104](#).

Parameters

Use the following parameters to customize your report:

- NetApp ONTAP Storage
- Volume
- Number of Largest Files to View
- Export Date Format

Set the date format to use when exporting data.

The default report parameters report on the ten largest files on all storage systems and volumes. Select up to 100 files.

Quick view

The Quick View section displays a pie chart of the largest files on your storage systems compared to the total storage. Use the NetApp ONTAP Storage parameter to display volumes on all storage systems or a specific storage system.

Note: The Quick View section is modified through the Number of Largest Files to View, NetApp ONTAP Storage, and Volume parameters.

Detail view

The following fields and corresponding data display in the Detail View section of the Files by Size report:

Node

The physical server where your files are stored.

Volume

The name of the volume on the node.

File

The name of the file returned by the report, including the path.

Owner

The assigned volume owner or node where the volume resides. To view file ownership information, create and run a NetApp ONTAP File Inventory job with the Traversal Mode set to Filewalk in conjunction with the IBM Storage Copy Data Management Filewalker tool before running this report

No. of Copies

The number of copies of the file that exist in the Inventory.

SnapMirror

The SnapMirror associated with the file.

SnapVault

The SnapVault associated with the file.

Size

The amount of space on the node used by the file.

Last Time Modified

The date and time in which the file was last modified.

Related concepts

[“File analytics reports” on page 243](#)

The File Analytics Reports help you review your storage needs and examine your storage capacity. There are a number of variables that you can use to view the files on your system. Using the File Analytics Reports you can drill down to information about the size of files, the age of files, and type of files on your storage systems. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

Protection compliance reports

The Protection Compliance reports help ensure your data is protected through user-defined recovery point objective parameters.

Reports are based on the data collected by the most recently run job. You must catalog all volumes on a storage system using full storage system cataloging to view the correct storage protection status for all volumes and qtrees.

Use the Protection Compliance reports to answer questions such as:

- Which of my NetApp ONTAP or VMware protection jobs have never run successfully?
- What is the remaining compliance time of a specific NetApp ONTAP storage system or VMware object?

Choose the Protection Compliance report that fits your needs.

- **File System RPO Compliance Report** - Displays the RPO compliance of file systems.
- **IBM Storage Virtualize RPO Compliance Report** - Displays IBM Storage Virtualize storage systems in relation to your recovery point objective parameters
- **NetApp ONTAP Protection Usage Report** - Displays the storage usage of IBM Storage Copy Data Management protection jobs on the volumes of your NetApp ONTAP storage systems.
- **NetApp ONTAP RPO Compliance Report** - Displays the primary snapshot protection for NetApp ONTAP storage systems.
- **Pure Storage FlashArray RPO Compliance** - Displays Pure Storage FlashArray storage systems in relation to your recovery point objective parameters.
- **Recovery Points Report** - Displays recovery points of various protected resources.
- **Unprotected Virtual Machines Report** - Displays the virtual machines that are not protected as a part of primary snapshot protection.
- **VMware RPO Compliance Report** - Displays the primary snapshot protection for virtual machines and datastores.

Related concepts

[“File System RPO compliance report” on page 252](#)

The File System RPO Compliance report displays your file systems in relation to your recovery point objective parameters. Determine which of your file systems are not in compliance with your RPO parameters, and discover the reasons for their non-compliance.

[“IBM Storage Virtualize RPO compliance report” on page 254](#)

The IBM Storage 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.

[“NetApp ONTAP protection usage report” on page 257](#)

The NetApp ONTAP Protection Space Usage report displays the storage usage of IBM Storage Copy Data Management protection jobs. Review the amount of space occupied by snapshots, SnapVaults, and SnapMirrors on the volumes of your NetApp ONTAP storage systems.

[“NetApp ONTAP RPO compliance report” on page 259](#)

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.

[“Pure Storage FlashArray RPO compliance report” on page 261](#)

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.

[“Recovery points report” on page 264](#)

The Recovery Points report displays recovery points of protected resources. View the protection type, job name, and location of your available recovery points.

[“Unprotected Virtual Machines report” on page 265](#)

The Unprotected Virtual Machines report displays virtual machines that are not protected with a primary snapshot or VM Replication. Quickly view all of your unprotected virtual machines, along with their eligibility for protection.

[“VMware RPO compliance report” on page 266](#)

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.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database

condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

[“Downloading a Report” on page 229](#)

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.

File System RPO compliance report

The File System RPO Compliance report displays your file systems in relation to your recovery point objective parameters. Determine which of your file systems are not in compliance with your RPO parameters, and discover the reasons for their non-compliance.

BEFORE YOU BEGIN:

- Create and run a File System Inventory job. You can select one or more file system in a single job definition for cataloging. See [“Creating a File System Inventory job definition” on page 99](#).
- Create and run a File System Backup job. See [“Creating a File System Backup job” on page 129](#).

Parameters

Use the following parameters to customize your report:

- Host

Multiple selections are supported.

- Protection Type

Set the protection types to return in the report. Values include Primary and Replication. Multiple selects are supported.

- Storage Vendor

Set the storage vendor types to display in the report. Multiple selections are supported.

- Display File Systems That Are

Set the compliance status of your file system to return in the report. Values include Compliant and Not Compliant. Multiple selections are supported. By default, this parameter is set to Not Compliant.

- RPO Older Than

Set the age of the recovery point objective in days.

The default report parameters report on all non-compliant file systems based on an RPO older than one day.

Quick View

The Quick View section displays a bar graph of compliant and non-compliant file systems based on your RPO parameters.

Not Compliant for Primary Protection

The following fields and corresponding data display in the Not Compliant for Primary Protection section of the File System RPO Compliance report:

File System Mount Point(s)

The file system's associated mount points.

Host (OS Type)

The hostname and operating system type. For example, Windows, AIX, or Linux.

Job Name

The Backup job associated with the file system

Storage Vendor

The storage vendor associated with the file system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Reason

The reason the file system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Primary Protection

The following fields and corresponding data display in the Compliant for Primary Protection section of the File System RPO Compliance report:

File System Mount Point(s)

The file system's associated mount points.

Host (OS Type)

The hostname and operating system type. For example, Windows, AIX, or Linux.

Job Name

The Backup job associated with the file system.

Storage Vendor

The storage vendor associated with the file system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your file system will be non-compliant.

Not Compliant for Replication

The following fields and corresponding data display in the Not Compliant for Replication section of the File System RPO Compliance report:

File System Mount Point(s)

The file system's associated mount points.

Source

The source of the secondary protection.

Destination

The destination of the secondary protection.

Job Name

The Backup job associated with the file system.

Protection Time

The most recent instance of a successful run of the Backup job.

Storage Vendor

The storage vendor associated with the file system.

Reason

The reason the file system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk..

Compliant for Replication

The following fields and corresponding data display in the Compliant for Replication section of the File System RPO Compliance report:

File System Mount Point(s)

The file system's associated mount points.

Host (OS Type)

The hostname and operating system type. For example, Windows, AIX, or Linux.

Job Name

The Backup job associated with the file system.

Storage Vendor

The storage vendor associated with the file system.

Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your file system will be non-compliant.

Related concepts

[“Protection compliance reports” on page 250](#)

The Protection Compliance reports help ensure your data is protected through user-defined recovery point objective parameters.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

IBM Storage Virtualize RPO compliance report

The IBM Storage 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 Storage Virtualize Inventory job. You can select one or more IBM provider in a single job definition for cataloging. See [“Creating an IBM Storage Virtualize Inventory job definition” on page 100](#).
- Create and run an IBM Storage Virtualize Backup job. See [“Creating an IBM Storage Virtualize Backup job definition” on page 131](#).

Use the IBM Storage Virtualize RPO Compliance report to answer questions such as:

- Which of my IBM storage systems are not RPO compliant for Flash Copy protection?
- Which of my IBM Backup jobs have never run successfully?

Parameters

Use the following parameters to customize your report:

- Storage Array

Multiple selections are supported.

- Protection Type

Set the IBM protection type to return in the report. Values include FlashCopy and Global Mirror with Change Volumes. Multiple selections are supported.

- Display Resources That Are

Set the compliance status of your IBM storage systems to return in the report. Values include Compliant and Not Compliant. Multiple selections are supported. By default, this parameter is set to Not Compliant.

- RPO Older Than

Set the age of the recovery point objective in days.

The default report parameters report on all non-compliant IBM storage systems based on an RPO older than one day.

Quick View

The Quick View section displays a bar graph of compliant and non-compliant IBM storage systems based on your RPO parameters.

Not Compliant for Flash Copy Protection

The following fields and corresponding data display in the Not Compliant for Flash Copy Protection section of the IBM Storage Virtualize RPO Compliance report:

Volume

The name of the IBM storage system.

Location

The location of the IBM storage system.

Consistency Group

The name of the volume's associated consistency group.

Job Name

The Backup job associated with the IBM storage system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Reason

The reason the IBM storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Flash Copy Protection

The following fields and corresponding data display in the Compliant for Snapshot Protection section of the IBM Storage Virtualize RPO Compliance report:

Volume

The name of the IBM storage system.

Location

The location of the IBM storage system.

Consistency Group

The name of the volume's associated consistency group.

Job Name

The Backup job associated with the IBM storage system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your IBM storage system will be non-compliant. Hover over the bar to see the remaining compliance time.

Not Compliant for Global Mirror with Change Volumes Protection

The following fields and corresponding data display in the Not Compliant for Global Mirror with Change Volumes Protection section of the IBM Storage Virtualize RPO Compliance report:

Volume

The name of the primary source volume.

Source

The source of the secondary protection.

Consistency Group

The name of the volume's associated consistency group.

Destination

The destination of the secondary protection.

Job Name

The Global Mirror job associated with the IBM storage system.

Protection Time

The most recent instance of a successful run of the protection job.

Reason

The reason the IBM storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Global Mirror with Change Volumes Protection

The following fields and corresponding data display in the Compliant for Global Mirror with Change Volumes Protection section of the IBM Storage Virtualize RPO Compliance report:

Volume

The name of the primary source volume.

Source

The source of the secondary protection.

Consistency Group

The name of the volume's associated consistency group.

Destination

The destination of the secondary protection.

Job Name

The Backup job associated with the IBM storage system.

Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your IBM storage system will be non-compliant. Hover over the bar to see the remaining compliance time.

Related concepts

[“Protection compliance reports” on page 250](#)

The Protection Compliance reports help ensure your data is protected through user-defined recovery point objective parameters.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

NetApp ONTAP protection usage report

The NetApp ONTAP Protection Space Usage report displays the storage usage of IBM Storage Copy Data Management protection jobs. Review the amount of space occupied by snapshots, SnapVaults, and SnapMirrors on the volumes of your NetApp ONTAP storage systems.

BEFORE YOU BEGIN:

- Create and run a NetApp ONTAP Storage Inventory job. You can select one or more NetApp ONTAP cluster providers in a single job definition for cataloging. See [“Creating a NetApp ONTAP Storage Inventory job definition” on page 103](#).
- Create and run a NetApp ONTAP Backup job. See [“Creating a NetApp ONTAP Backup job definition” on page 140](#).

Use the NetApp ONTAP Protection Space Usage report to answer questions such as:

- What is the secondary protection storage usage across my volumes?
- What is the combined size of all of my volumes on a NetApp ONTAP storage system?

Parameters

Use the following parameters to customize your report:

- Storage Array
 - Multiple selections are supported.
- Protection Type
 - Set the NetApp ONTAP protection type. Protection types include SnapMirror, SnapVault, and Snapshot. Multiple selections are supported.
- Job Type
 - Set the IBM Storage Copy Data Management protection jobs to display in the report. Multiple selections are supported.

Quick View

The Quick View section displays a bar graph of the storage protection usage across all of the volumes on the selected NetApp ONTAP storage system.

Storage Used by Snapshot Protection

The following fields and corresponding data display in the Storage Used by Snapshot Protection section of the NetApp ONTAP Protection Usage report:

Primary Source Volume

The name of the primary source volume.

Location

The location of the primary source volume.

Job Names

The names of the associated IBM Storage Copy Data Management protection jobs.

Snapshot Count

The number of snapshots available on the volume.

Oldest Snapshot Creation Time

The creation date and time of the oldest snapshot on the volume.

Volume Size

The total size of the volume.

Volume Used Size

The size of the volume occupied by data.

Snapshot Usage

The amount of space on the volume dedicated to snapshot protection.

Storage Used by SnapVault/SnapMirror Protection

The following fields and corresponding data display in the Storage Used by SnapVault/SnapMirror Protection section of the NetApp ONTAP Protection Usage report:

Primary Source Volume

The name of the primary source volume.

Destination

The destination of the secondary protection.

Job Names

The names of the associated IBM Storage Copy Data Management protection jobs.

SnapVault Usage

The amount of space on the volume dedicated to SnapVault protection.

SnapMirror Usage

The amount of space on the volume dedicated to SnapMirror protection.

Related concepts

[“Protection compliance reports” on page 250](#)

The Protection Compliance reports help ensure your data is protected through user-defined recovery point objective parameters.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

NetApp ONTAP RPO compliance report

The NetApp ONTAP RPO Compliance report displays NetApp ONTAP storage systems in relation to your recovery point objective parameters. Determine which of your NetApp ONTAP storage systems are not in compliance with your RPO parameters, and discover the reasons for their non-compliance.

BEFORE YOU BEGIN:

- Create and run a NetApp ONTAP Storage Inventory job. You can select one or more NetApp ONTAP cluster providers in a single job definition for cataloging. See [“Creating a NetApp ONTAP Storage Inventory job definition” on page 103](#).
- Create and run a NetApp ONTAP Backup job. See [“Creating a NetApp ONTAP Backup job definition” on page 140](#).

Use the NetApp ONTAP RPO Compliance report to answer questions such as:

- Which of my NetApp ONTAP storage systems are not RPO compliant for SnapVault or SnapMirror protection?
- Which of my NetApp ONTAP Backup jobs have never run successfully?

Parameters

Use the following parameters to customize your report:

- Storage Array
Multiple selections are supported.
- Protection Type
Set the RPO compliance protection type. Protection types include SnapMirror, SnapVault, and Snapshot. Multiple selections are supported. By default, this parameter is set to All.
- Display Resources That Are
Set the compliance status of your NetApp ONTAP storage systems to return in the report. Values include Compliant and Not Compliant. Multiple selections are supported. By default, this parameter is set to Not Compliant.
- RPO Older Than
Set the age of the recovery point objective in days.

The default report parameters report on all non-compliant NetApp ONTAP storage systems based on an RPO older than one day.

Quick View

The Quick View section displays a bar graph of compliant and non-compliant NetApp ONTAP storage systems based on your RPO parameters.

Not Compliant for Primary Protection

The following fields and corresponding data display in the Not Compliant for Primary Protection section of the NetApp ONTAP RPO Compliance report:

Resource Name

The name of the NetApp ONTAP storage system.

Location

The location of the NetApp ONTAP storage system.

Job Name

The Backup job associated with the NetApp ONTAP storage system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Reason

The reason the NetApp ONTAP storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Primary Protection

The following fields and corresponding data display in the Compliant for Primary Protection section of the NetApp ONTAP RPO Compliance report:

Resource Name

The name of the NetApp ONTAP storage system.

Location

The location of the NetApp ONTAP storage system.

Job Name

The Backup job associated with the NetApp ONTAP storage system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your NetApp ONTAP storage system will be non-compliant. Hover over the bar to see the remaining compliance time.

Not Compliant for Replication

The following fields and corresponding data display in the Not Compliant for Replication section of the NetApp ONTAP RPO Compliance report:

Source

The source of the secondary protection.

Destination

The destination of the secondary protection.

Job Name

The SnapVault or SnapMirror Backup job associated with the NetApp ONTAP storage system.

Protection Time

The most recent instance of a successful run of the protection job.

Reason

The reason the NetApp ONTAP storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Replication

The following fields and corresponding data display in the Compliant for Replication section of the NetApp ONTAP RPO Compliance report:

Source

The source of the secondary protection.

Destination

The destination of the secondary protection.

Job Name

The Backup job associated with the NetApp ONTAP storage system.

Protection Time

The most recent instance of a successful run of the Backup job.

Remaining Compliance Time

The time remaining before your NetApp ONTAP storage system will be non-compliant. Hover over the bar to see the remaining compliance time.

Related concepts

[“Protection compliance reports” on page 250](#)

The Protection Compliance reports help ensure your data is protected through user-defined recovery point objective parameters.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

Pure Storage FlashArray RPO compliance report

The Pure Storage FlashArray RPO Compliance report displays Pure Storage systems in relation to your recovery point objective parameters. Determine which of your Pure Storage systems are not in compliance with your RPO parameters, and discover the reasons for their non-compliance.

BEFORE YOU BEGIN:

- Create and run a Pure Storage Inventory job. You can select one or more Pure Storage provider in a single job definition for cataloging. See [“Creating a Pure Storage FlashArray Inventory job definition” on page 107](#).
- Create and run a Pure Storage Backup job. See [“Creating a Pure Storage FlashArray Backup job definition” on page 142](#).

Parameters

Use the following parameters to customize your report:

- Storage Array

Multiple selections are supported.

- Protection Type

Set the Pure Storage FlashArray protection type to return in the report. Values include Primary and Replication. Multiple selections are supported.

- Display Resources That Are

Set the compliance status of your Pure Storage systems to return in the report. Values include Compliant and Not Compliant. Multiple selections are supported. By default, this parameter is set to Not Compliant.

- RPO Older Than

Set the age of the recovery point objective in days.

Quick View

The Quick View section displays a bar graph of compliant and non-compliant Pure Storage FlashArray systems based on your RPO parameters.

Not Compliant for Primary Protection

The following fields and corresponding data display in the Not Compliant for Primary Protection section of the Pure Storage FlashArray RPO Compliance report:

Volume

The name of the Pure Storage FlashArray system volume.

Location

The location of the Pure Storage FlashArray system.

Job Name

The Backup job associated with the Pure Storage FlashArray system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Reason

The reason the Pure Storage FlashArray system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Primary Protection

The following fields and corresponding data display in the Compliant for Primary Protection section of the Pure Storage FlashArray RPO Compliance report:

Volume

The name of the Pure Storage FlashArray system volume.

Location

The location of the Pure Storage FlashArray system.

Job Name

The Backup job associated with the Pure Storage FlashArray system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your Pure Storage FlashArray system will be non-compliant.

Not Compliant for Replication

The following fields and corresponding data display in the Not Compliant for Replication section of the Pure Storage FlashArray RPO Compliance report:

Volume

The name of the Pure Storage FlashArray system volume.

Source Storage Array

The source of the secondary protection.

Destination Storage Array

The destination of the secondary protection.

Job Name

The Backup job associated with the Pure Storage FlashArray storage system.

Last Successful Protection Time

The most recent instance of a successful run of the protection job.

Reason

The reason the Pure Storage FlashArray storage system does not meet your RPO compliance parameters. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Replication

The following fields and corresponding data display in the Compliant for Replication section of the Pure Storage FlashArray RPO Compliance report:

Volume

The name of the Pure Storage FlashArray system volume.

Source Storage Array

The source of the secondary protection.

Destination Storage Array

The destination of the secondary protection.

Job Name

The Backup job associated with the Pure Storage FlashArray storage system.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Compliance Time Remaining

The time remaining before your Pure Storage FlashArray storage system will be non-compliant. Hover over the bar to see the remaining compliance time.

Related concepts

[“Protection compliance reports” on page 250](#)

The Protection Compliance reports help ensure your data is protected through user-defined recovery point objective parameters.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

Recovery points report

The Recovery Points report displays recovery points of protected resources. View the protection type, job name, and location of your available recovery points.

BEFORE YOU BEGIN:

- Create and run a Backup job.

Parameters

Use the following parameters to customize your report:

- Site

Multiple selections are supported.

- Resource Type

Set the resource type to return in the report. Resource types include Oracle Database, SAP HANA Database, SQL Database, VMware Datastores, and VMware VMs. By default, this parameter is set to All.

- Protection Type

Set the protection type to display in the report. Protection types include primary and secondary. Multiple selections are supported. By default, this parameter is set to All.

- Show Recovery Points

Set the age of the recovery point objective in days.

Recovery Points

The following fields and corresponding data display in the Recovery Points section of the Recovery Points report:

Resource Name

The name and location of the protected resource.

Site

The site associated with the protected resource.

Protection Time

The most recent instance of a successful run of the Backup job.

Protection Type

The recovery point's protection type.

Job Name

The Backup job associated with the recovery point.

Volumes (Location)

The volumes associated with the displayed recovery point.

Related concepts

[“Protection compliance reports” on page 250](#)

The Protection Compliance reports help ensure your data is protected through user-defined recovery point objective parameters.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

Unprotected Virtual Machines report

The Unprotected Virtual Machines report displays virtual machines that are not protected with a primary snapshot or VM Replication. Quickly view all of your unprotected virtual machines, along with their eligibility for protection.

BEFORE YOU BEGIN:

- Create and run a VMware Inventory job. See [“Creating a VMware Inventory job definition” on page 109](#).
- All related storage providers must be added to IBM Storage Copy Data Management, which include NetApp storage controllers and clusters. Create and run a NetApp ONTAP Storage Inventory job to catalog associated NetApp ONTAP storage providers. See [“Creating a NetApp ONTAP Storage Inventory job definition” on page 103](#).
- Create and run a VMware Backup job. See [“Creating a VMware Backup job definition” on page 145](#).

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.

Datstores

The name of the associated datastore.

Related concepts

[“Protection compliance reports” on page 250](#)

The Protection Compliance reports help ensure your data is protected through user-defined recovery point objective parameters.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

VMware RPO compliance report

The VMware RPO Compliance report displays VMware objects in relation to your recovery point objective parameters. Determine which of your VMware objects are not in compliance with your RPO parameters, and discover the reasons for their non-compliance.

BEFORE YOU BEGIN:

- Create and run a VMware Inventory job. See [“Creating a VMware Inventory job definition” on page 109](#).
- All related NetApp ONTAP storage providers must be added to IBM Storage Copy Data Management, which include NetApp ONTAP storage controllers and clusters. Create and run a NetApp ONTAP Storage Inventory job to catalog associated NetApp ONTAP storage providers. See [“Creating a NetApp ONTAP Storage Inventory job definition” on page 103](#).
- Create and run a VMware Backup job. See [“Creating a VMware Backup job definition” on page 145](#).

Use the VMware RPO Compliance report to answer questions such as:

- Which of my VMware Backup job have never run successfully?
- What is the remaining compliance time of a specific VMware object?

Parameters

Use the following parameters to customize your report:

- vCenter

Multiple selections are supported.

- Primary Protection Resource

Set the primary protection resource type to display in the report. Resource types include Datastore and Virtual Machine. Multiple selections are supported. By default, this parameter is set to All.

- Replication Resource

Set the replication resource type to display in the report. Resource types include Datastore and Virtual Machine. Multiple selections are supported. By default, this parameter is set to All.

- Protection Type

Set the RPO compliance protection type. Protection types include Primary and Replication. Multiple selections are supported. By default, this parameter is set to All.

- Storage Vendor

Set the storage vendor types to display in the report. Multiple selections are supported.

- Display Resources That Are

Set the compliance status of your VMware objects to return in the report. Values include Compliant and Not Compliant. Multiple selections are supported. By default, this parameter is set to Not Compliant.

- RPO Older Than

Set the age of the recovery point objective in days.

- View VMs with Application(s) only

Select Yes to view only application virtual machines. Select the application type through the Application Server parameter. Select No to display all virtual machines.

The default report parameters report on all non-compliant VMware objects based on an RPO older than one day.

Quick View

The Quick View section displays a bar graph of compliant and non-compliant VMware objects based on your RPO parameters.

Not Compliant for Primary Protection

The following fields and corresponding data display in the Not Compliant for Primary Protection section of the VMware RPO Compliance report:

Resource Name

The name of the virtual machine, preceded by an icon that displays the object type or a virtual machine's power state.

Location

The location of the VMware object.

Job Name

The Backup job associated with the VMware object.

Storage Vendor

The storage vendor associated with the VMware object.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Application(s)

The application server type, if applicable. Available types include Active Directory Server, Exchange Server, SQL Server, and SharePoint Server.

Reason

The reason the VMware object is not compliant for primary protection. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Primary Protection

The following fields and corresponding data display in the Compliant for Primary Protection section of the VMware RPO Compliance report:

Resource Name

The name of the object, preceded by an icon that displays the object type or a virtual machine's power state.

Location

The location of the VMware object.

Job Name

The Backup job associated with the VMware object.

Storage Vendor

The storage vendor associated with the VMware object.

Last Successful Protection Time

The most recent instance of a successful run of the Backup job.

Application(s)

The application server type, if applicable. Available types include Active Directory Server, Exchange Server, SQL Server, and SharePoint Server.

Compliance Time Remaining

The time remaining before your VMware object will be non-compliant.

Not Compliant for Replication

The following fields and corresponding data display in the Not Compliant for Replication section of the VMware RPO Compliance report:

Resource Name

The name of the virtual machine.

Source

The source of the secondary protection.

Destination

The destination of the secondary protection.

Job Name

The Backup job associated with the VMware object.

Protection Time

The most recent instance of a successful run of the Backup job.

Storage Vendor

The storage vendor associated with the VMware object.

Reason

The reason the VMware object is not compliant for secondary protection. Examples include no successful runs of a protection job, or backing up to an unsupported disk.

Compliant for Replication

The following fields and corresponding data display in the Compliant for Replication section of the VMware RPO Compliance report:

Resource Name

The name of the virtual machine.

Source

The source of the secondary protection.

Destination

The destination of the secondary protection.

Job Name

The Backup job associated with the VMware object.

Protection Time

The most recent instance of a successful run of the Backup job.

Storage Vendor

The storage vendor associated with the VMware object.

Compliance Time Remaining

The time remaining before your VMware object will be non-compliant.

Related concepts

[“Protection compliance reports” on page 250](#)

The Protection Compliance reports help ensure your data is protected through user-defined recovery point objective parameters.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

Storage protection reports

The Storage Protection Reports help ensure your data is protected and display the status of your replication process. Use the Storage Protection Reports to view qtrees that are protected using NetApp ONTAP SnapVault or NetApp ONTAP SnapMirror functionality and volumes that are protected using NetApp ONTAP SnapMirror functionality. You can also view unprotected qtrees and volumes to help plan your data protection strategies.

Reports are based on the data collected by the most recently run job. You must catalog all volumes on a storage system using full storage system cataloging to view the correct storage protection status for all volumes and qtrees.

Use the Storage Protection Reports to answer questions such as:

- What 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?

Note: IBM Storage 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.

Choose the Storage Protection report that fits your needs.

- **NetApp ONTAP OSSV Relationship Status Report**- Review the relationship status of your OSSV protected clients.
- **NetApp ONTAP Overprotected Volumes** - View volumes that are overprotected using NetApp ONTAP SnapVault or SnapMirror functionality.
- **NetApp ONTAP Qtree Protection Status Report** - View qtrees that are protected using NetApp ONTAP SnapVault or SnapMirror functionality.
- **NetApp ONTAP Underprotected Volumes** - View volumes that are underprotected using NetApp ONTAP SnapVault, SnapMirror, or Snapshot functionality.

- **NetApp ONTAP Volume Protection Status Report** - View volumes that are protected using NetApp ONTAP SnapVault, SnapMirror, or Snapshot functionality and evaluate your volume replication processes.
- **NetApp ONTAP Transition Dependency** - View inter-node replication dependencies between sources and destinations for 7-mode systems.

Related concepts

[“NetApp ONTAP OSSV relationship status report” on page 271](#)

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.

[“NetApp ONTAP overprotected volumes report” on page 272](#)

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.

[“NetApp ONTAP Qtree protection status report” on page 273](#)

The NetApp ONTAP Qtree Protection Status report displays qtrees that are protected using NetApp ONTAP Snapshot, SnapVault, or SnapMirror software.

[“NetApp ONTAP underprotected volumes report” on page 274](#)

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.

[“NetApp ONTAP volume protection status report” on page 275](#)

The NetApp ONTAP Volume Protection Status report displays volumes that are protected or unprotected using NetApp ONTAP Snapshot, SnapVault, or SnapMirror software.

[“NetApp ONTAP transition dependency report” on page 277](#)

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.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

[“Downloading a Report” on page 229](#)

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.

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 concepts

[“Storage protection reports” on page 269](#)

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.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

NetApp ONTAP overprotected volumes report

The NetApp ONTAP Overprotected Volumes report displays volumes that are overprotected using NetApp ONTAP Snapshot, SnapVault, or SnapMirror software. Through your parameters selections, you can decide what constitutes an overprotected volume. Select the number of acceptable snapshots to reside on a volume, and if a volume should also be SnapMirrored and SnapVaulted. If a volume exceeds any of your parameter selections, it will be returned on the NetApp ONTAP Overprotected Volumes report.

Note: This report may discover storage systems that have not been cataloged in the database. These non-cataloged storage systems are discovered due to their replication relationships with other cataloged storage systems.

Use the NetApp ONTAP Overprotected Volumes report to answer questions such as:

- How many of my volumes are overprotected?
- How many gigabytes of space are dedicated to overprotecting a volume?
- Which of my volumes are protected with more than 10 snapshots?

Parameters

Use the following parameters to customize your report:

- Storage Array
Multiple selections are supported.
- Acceptable Snapshots
- Is Volume SnapMirrored?
- Is Volume SnapVaulted?

The default report parameters display all overprotected volumes with ten or more snapshots.

Detail View - Overprotected Volumes

The following fields and corresponding data display in the Detail View - Overprotected Volumes section of the NetApp ONTAP Overprotected Volumes report:

Volume

The name of the overprotected volume.

Storage Array

The host where the overprotected volume is located.

Location

The node where the volume resides.

Overprotection Reason

The reason the volume was returned by the report. For example, if the number of snapshots is larger than the defined Acceptable Snapshots parameter, or if a volume is SnapVaulted and the Is Volume SnapVaulted parameter is set to No.

Overprotection Storage Cost

The amount of space on the volume dedicated to overprotection.

Related concepts

[“Storage protection reports” on page 269](#)

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.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

NetApp ONTAP Qtree protection status report

The NetApp ONTAP Qtree Protection Status report displays qtrees that are protected using NetApp ONTAP Snapshot, SnapVault, or SnapMirror software.

Note: This report may discover storage systems that have not been cataloged in the database. These non-cataloged storage systems are discovered due to their replication relationships with other cataloged storage systems.

Use the NetApp ONTAP Qtree Protection Status report to answer questions such as:

- Which of my cataloged SnapVaulted or SnapMirrored Qtrees are in an unprotected state?
- Which of my NetApp ONTAP volumes are exceeding their lag time by 30 days or more?

Parameters

Use the following parameters to customize your report:

- Storage Array
 - Multiple selections are supported.
- Protection Type
- Lag time in days
- Show Protected Qtrees
- Record Limit

The default report parameters display the unprotected qtrees for all storage systems and all volumes.

Detail View

The following fields and corresponding data display in the following sections: NetApp ONTAP SnapMirrored Qtrees in an Unprotected State, Qtrees Protected by SnapMirror, NetApp ONTAP SnapVaulted Qtrees in an Unprotected State, and Qtrees Protected by SnapVault.

Storage Array

The physical server where your files are stored.

Source

The name of the protected or unprotected qtree.

Destination

The node where the replication destination is located.

State

The state of the destination (for example, SnapMirrored, SnapVaulted, broken-off, uninitialized, or unknown).

Lag Time

The lag time in days, hours, minutes, and seconds between the source and the destination.

Xfer Throughput

The transfer throughput in KBs per second between the source and the destination.

Related concepts

[“Storage protection reports” on page 269](#)

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.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

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 parameter selections, you can decide what constitutes an underprotected volume. Select the number of acceptable snapshots to reside on a volume, the lag time, and if a volume should also be SnapMirrored and SnapVaulted. If a volume has less protection than your parameter selections, it will be returned on the NetApp ONTAP Underprotected Volumes report.

Note: This report may discover storage systems that have not been cataloged in the database. These non-cataloged storage systems are discovered due to their replication relationships with other cataloged storage systems.

Use the NetApp ONTAP Underprotected Volumes report to answer questions such as:

- What are the names of my underprotected volumes?
- Which of my volumes are underprotected with less than 10 snapshots?

Parameters

Use the following parameters to customize your report:

- Storage Array

Multiple selections are supported.

- Acceptable Snapshots
- Lag time in days
- Is Volume SnapMirrored?
- Is Volume SnapVaulted?

The default report parameters display underprotected volumes with less than ten snapshots that are older than 7 days.

Underprotected Volumes

The following fields and corresponding data display in the Underprotected Volumes section of the NetApp ONTAP Underprotected Volumes report:

Volume

The name of the underprotected volume.

Storage Array

The node where the underprotected volume is located.

Location

The node where the volume resides.

Underprotection Reason

The reason the volume was returned by the report. For example, if the number of snapshots is lower than the defined No. of Acceptable Snapshots parameter, or if the snapshots have excessive lag times.

Related concepts

[“Storage protection reports” on page 269](#)

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.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

NetApp ONTAP volume protection status report

The NetApp ONTAP Volume Protection Status report displays volumes that are protected or unprotected using NetApp ONTAP Snapshot, SnapVault, or SnapMirror software.

Note: This report may discover storage systems that have not been cataloged in the database. These non-cataloged storage systems are discovered due to their replication relationships with other cataloged storage systems.

Use the NetApp ONTAP Volume Protection Status report to answer questions such as:

- Which of my cataloged SnapVaulted or SnapMirrored volumes are in an unprotected state?
- Which of my NetApp ONTAP volumes are exceeding their lag time by 30 days or more?

Parameters

Use the following parameters to customize your report:

- Storage Array
Multiple selections are supported.
- Protection Type
- Lag time in days
- Show Protected Volumes

The default report parameters display all unprotected SnapMirror volumes.

Detail View - NetApp ONTAP SnapMirrored and SnapVaulted Volumes

The following fields and corresponding data display in the following sections: NetApp ONTAP SnapMirrored Volumes in an Unprotected State, NetApp ONTAP Volumes Protected by SnapMirror, NetApp ONTAP SnapVaulted Volumes in an Unprotected State and NetApp ONTAP Volumes Protected by SnapVault.

Storage Array

The node on which the protected or unprotected volume is located.

Source

The name of the protected or unprotected volume.

Destination

The node on which the replication destination is located.

State

The state of the destination. For example, SnapMirrored, SnapVaulted, broken-off, uninitialized, or unknown.

Type

The SnapMirror type for clustered volumes. For example, Vault or Mirror. Blank entries indicate non-clustered 7-Mode volumes.

Lag Time

The lag time in days, hours, minutes, and seconds between the source and the destination.

Xfer Throughput (NetApp ONTAP Volumes Protected by SnapMirror/SnapVault section only)

The transfer throughput in KBs per second between the source and the destination.

Detail View - NetApp ONTAP Snapshotted Volumes

The following fields and corresponding data display in the following sections NetApp ONTAP Volume with Snapshots Exceeding Lag Time and NetApp ONTAP Volumes Protected by Snapshot.

Storage Array

The node on which the protected or unprotected volume is located.

Volume

The name of the protected or unprotected volume.

Location

The node where the volume resides.

Snapshot Count

The number of snapshots available on the volume.

Latest Snapshot Time

The date and time of the most recent snapshot on the volume.

Related concepts

[“Storage protection reports” on page 269](#)

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.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

NetApp ONTAP transition dependency report

The NetApp ONTAP Transition Dependency report displays the destination nodes and replication relationships for all of your NetApp ONTAP storage systems operating in 7-mode.

Before you begin

- Create and run a NetApp ONTAP Storage Inventory job to catalog your NetApp ONTAP storage systems operating in 7-mode. You can select one or more NetApp ONTAP cluster provider in a single job definition for cataloging. See [“Creating a NetApp ONTAP Storage Inventory job definition” on page 103](#).
- Create and run a NetApp ONTAP Backup job. See [“Creating a NetApp ONTAP Backup job definition” on page 140](#).

Parameters

Use the following parameters to customize your report:

- Storage Array
Set the NetApp ONTAP storage systems to display in the report. Multiple selections are supported.
- Show Detailed View
Enable to show detailed volume, SnapMirror, and SnapVault data in the report.

Summary View

The following fields and corresponding data display in the Summary View section of the NetApp ONTAP Transition Dependency report:

Storage Array

The name of the NetApp ONTAP storage system source, along with the operating system version running on the source node.

SnapMirror Destination Nodes (Count)

The name of the SnapMirror destination node. The relationship count displays in parentheses.

SnapVault Destination Nodes (Count)

The name of the SnapVault destination node. The relationship count displays in parentheses.

Destination Node OS Version

The operating system version running on the destination node.

Detail View

If the Show Detailed View parameter is set to Yes, the following fields and corresponding data display in the Detail View section of the NetApp ONTAP Transition Dependency report:

Storage Array: Volume

The name of the NetApp ONTAP storage system and associated source volume, along with the size of the source volume.

SnapMirror Destination

The name of the SnapMirror destination and associated volume.

SnapVault Destination

The name of the SnapVault destination and associated volume.

Related concepts

[“Storage protection reports” on page 269](#)

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.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

Storage utilization reports

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

Use the Storage Utilization Reports to answer questions such as:

- What is the total available storage space in the entire system?
- What is the amount of free and used space on my volumes?
- How many files and disks make up my aggregates?

Note: Utilization reports provide advanced warning for volumes, aggregates, or LUNs that are beyond a specified capacity range. Similar reports are available for VMware datastores.

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.

- **IBM Storage Virtualize Consistency Groups Report** - Display information about your IBM Consistency Groups. View the associated storage providers, source and target volumes, and protection status of your IBM volumes through Consistency Groups
- **IBM Storage 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 Storage 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.
- **Instant Disk Restore Volumes Report** - Display a list of mapped volumes created through Instant Disk Restore jobs.
- **NetApp ONTAP Aggregates Report** - Review the storage utilization and configuration of your NetApp aggregates. View the total and free space available and the number of volumes and disks that make up your aggregates.
- **NetApp ONTAP LUNs Report** - Review the total capacity of your LUNs, the total free space, and the percentage available to ascertain your NetApp ONTAP LUN storage utilization.
- **NetApp ONTAP Orphaned LUNs Report** - Review NetApp ONTAP storage orphaned LUNs. These are the LUNs that have no initiator group mapping or belong to volumes that are offline.
- **NetApp ONTAP Quotas Report** - Review quota status to determine which users or groups are approaching or exceeding quota limits.
- **NetApp ONTAP Snapshots Report** - Review the total capacity of your snapshots, the total free space, and the percentage available to ascertain your NetApp ONTAP Snapshot storage utilization.
- **NetApp ONTAP Volumes Report** - Review the total capacity of your volumes, the total free space, and the percentage available to ascertain your NetApp ONTAP volume storage utilization.
- **Pure Storage FlashArray Volumes Report** - Review the total capacity of your volumes, the total free space, and the percentage available to ascertain your Pure Storage FlashArray volume storage utilization.
- **Storage Capacity Report** - Report the storage capacity of your IBM Storage Virtualize pools, and NetApp ONTAP aggregates.
- **VM and Storage Mapping Report** - Report that maps VMs all the way to the physical storage from which the datastore is created.
- **VMware Datastores Report** - Review the total capacity of your datastores, the total free space, and the percentage available to ascertain your VMware datastore storage utilization.
- **VMware LUNs Report** - Displays information about VMware LUNs such as which ESX server it belongs to, its datastore, vendor, total, and allocated capacity.
- **VMware Orphaned Datastores Report** - Review the datastores that do not have any virtual machines assigned to them, or if virtual machines are assigned to the datastores, view the virtual machines that are in an inaccessible state.
- **VMware Orphaned LUNs Report** - Review VMware orphaned LUNs. These are the LUNs not used as datastores or RDMs.
- **VMware VM Snapshot Sprawl Report** - Displays information about virtual machines with aged and memory snapshots.
- **VMware VM Sprawl Report** - Displays storage utilization across virtual machines based on their power state and storage utilization across virtual machine templates.
- **VMware VM Storage Report** - Review your virtual machines and associated datastores.

Quick View

This area of the report is a graphical illustration of the report using pie charts. For example, the quick view of the NetApp ONTAP Storage Volumes report shows the total capacity of your volumes, the free space, and the used space.

Summary View

This area of the report displays a summary of the data returned in the report. For example, the summary view of the NetApp ONTAP Storage Aggregates report shows the total used, free, and reserved space on your aggregate.

Detail View

This area of the report is a table where each row details a storage system, its corresponding volume or aggregate, and details returned by the report. For example, the NetApp ONTAP Storage Aggregates report shows the used and free space, volume count, disk count, and status of your aggregates.

Choose the Storage Utilization Report that fits your needs.

Related concepts

[“IBM Storage Virtualize Consistency Groups Report” on page 282](#)

Run the IBM Storage 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.

[“IBM Storage Virtualize Pools Report” on page 283](#)

Run the IBM Storage 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.

[“IBM Storage Virtualize Volumes Report” on page 285](#)

Report the storage utilization of your IBM storage volumes. Run the IBM Storage Virtualize Storage Volumes report to view the total capacity of your volumes, the total free space, and the available percentage.

[“Instant Disk Restore Volumes Report” on page 287](#)

Display a list of mapped volumes created through Instant Disk Restore jobs. View the name of the volume, location, and associated restore job.

[“NetApp ONTAP Aggregates Report” on page 288](#)

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.

[“NetApp ONTAP LUNs Report” on page 290](#)

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.

[“NetApp ONTAP Orphaned LUNs Report” on page 292](#)

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.

[“NetApp ONTAP Quotas Report” on page 293](#)

Review the quotas of your NetApp ONTAP storage devices to determine which users or groups are approaching or exceeding their quota limits.

[“NetApp ONTAP Snapshots Report” on page 295](#)

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.

[“NetApp ONTAP Volumes Report” on page 296](#)

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.

[“Pure Storage FlashArray Volumes Report” on page 298](#)

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.

[“Storage Capacity Report” on page 300](#)

Report the storage capacity of your IBM Storage Virtualize 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.

[“VM and storage mapping report” on page 301](#)

Review the mappings of virtual machines down to the physical storage from which the datastore is created.

[“VMware Datastores Report” on page 302](#)

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.

[“VMware LUNs Report” on page 304](#)

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.

[“VMware Orphaned Datastores Report” on page 306](#)

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” on page 307](#)

Review VMware orphaned LUNs. These are the LUNs that are neither used as datastores nor RDMs.

[“VMware VM Snapshot Sprawl Report” on page 309](#)

The VMware VM Snapshot Sprawl report displays the age and number of snapshots used to protect your virtual machines through NetApp ONTAP Snapshot software..

[“VMware VM Sprawl Report” on page 310](#)

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.

[“VMware VM Storage Report” on page 312](#)

Review your virtual machines and associated datastores through the VMware VM Storage report. View datastores, resource pools, and the provisioned space of the datastore.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

[“Downloading a Report” on page 229](#)

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.

IBM Storage Virtualize Consistency Groups Report

Run the IBM Storage Virtualize Consistency Group report to display information about your IBM Consistency Groups. View the associated storage providers, source and target volumes, and protection status of your IBM volumes through Consistency Groups.

Before you begin

Create and run an IBM Inventory job. You can select one or more IBM providers in a single job definition for cataloging. See [“Creating an IBM Storage Virtualize Inventory job definition” on page 100](#).

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 Storage Virtualize Consistency Groups report:

Consistency Group

The name of the Consistency Group along with its status (for example: copying or idle/copied) and FlashTime.

Storage Array

The name of the storage virtualizer as known to IBM Storage Copy Data Management where the consistency group resides.

Mapping Name

The mapping name IBM Storage 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 Storage Virtualize Consistency Groups report:

Consistency Group

The name of the consistency group along with its status (for example: copying or idle/copied) and FlashTime, master and auxiliary cluster.

Storage Array

The name of the storage virtualizer as known to IBM Storage Copy Data Management where the consistency group resides.

Mapping Name

The mapping name IBM Storage 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 concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

IBM Storage Virtualize Pools Report

Run the IBM Storage 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 Storage Virtualize Inventory job. You can select one or more IBM providers in a single job definition for cataloging. See [“Creating an IBM Storage Virtualize Inventory job definition” on page 100](#).

Use the IBM Storage 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 Storage Virtualize Pools report:

Storage Array

The name of the storage virtualizer as known to IBM Storage 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 Storage Virtualize Pools report:

Storage Pool

The name of the storage pool and associated child pools. Note that child pool capacities are not included in column totals.

Storage Array

The name of the storage virtualizer as known to IBM Storage Copy Data Management.

Capacity

The total amount of MDisk storage that is assigned to the storage pool.

Allocated

The total storage allocated to volumes within the storage pool.

Virtual Capacity

The total virtual size of all the volume copies that are associated with the storage pool.

External Virtual Capacity

The aggregate capacity of the managed and image mode MDisks from storage controllers virtualized using the “External Virtualization” feature of the chosen IBM storage systems.

Volumes

The number of volume copies that are in the storage pool.

Disk Count

The number of MDisks in the storage pool.

Status

The status of the MDisk with the highest priority status in the group, excluding image mode MDisks

Warning

A warning is generated when the assigned amount of space in the storage pool exceeds this level.

%

The percentage of space used on the storage pool.

% Used/Free

A status bar that displays the used space on the storage pool.

Related concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

IBM Storage Virtualize Volumes Report

Report the storage utilization of your IBM storage volumes. Run the IBM Storage 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 Storage Virtualize Inventory job. You can select one or more IBM providers in a single job definition for cataloging. See [“Creating an IBM Storage Virtualize Inventory job definition” on page 100](#).

Use the IBM Storage Virtualize Storage Volumes report to answer questions such as:

- What is the total available storage space in the entire system?
- What is the amount of free and used space?
- How many volumes are available on a specific storage system?
- What is the size of the volume and the storage system that it resides on?

Parameters

Use the following parameters to customize your report:

- Storage Array
- Detail View Filter
- Show Only Thin Provisioned Volumes

Quick View

The Quick View section displays the overall volume utilization of your IBM storage volumes.

Summary View

The following fields and corresponding data display in the Summary View section of the IBM Storage Virtualize Storage Volumes report:

Storage Array

The name of the IBM storage system.

Volume Count

The number of volumes available on the IBM storage system.

Capacity

The volume storage capacity that is available to a host.

Real Capacity

The amount of physical storage that is allocated from a storage pool to volume copies.

Used Capacity

The portion of real capacity that is being used to store data.

Free Capacity

The difference between the real capacity and used capacity values for volume copies.

Detail View

The following fields and corresponding data display in the Detail View section of the IBM Storage Virtualize Storage Volumes report:

Volume

The name of the volume on the IBM storage system.

Storage Array

The physical storage system where your files are stored.

Storage Pool

The volume's associated storage pool.

Capacity

The difference between the real capacity and used capacity values for volume copies.

Real Capacity

The amount of physical storage that is allocated from a storage pool to volume copies.

Used Capacity

The portion of real capacity that is being used to store data.

Warning Threshold

For thin provisioned or compressed volume copies, a warning is generated at this percentage of the volume capacity.

Status

The status of the volume. A volume can be online, offline, or degraded.

Thin

Displays the thin provisioned status of the volume.

% Used/Free

A status bar that displays the used space on the volume.

Related concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

Instant Disk Restore Volumes Report

Display a list of mapped volumes created through Instant Disk Restore jobs. View the name of the volume, location, and associated restore job.

Before you begin

Create and run an Instant Disk Restore job.

Parameters

Use the following parameters to customize your report:

- Storage Vendor

Set the storage vendor types to display in the report. Multiple selections are supported.

Detail View

The following fields and corresponding data display in the Detail View section of the Instant Disk Restore Volumes report:

Volume

The name of the mapped Instant Disk Restore volume on the node.

Location (Site)

The node and associated site where the volume resides.

Storage Vendor

The storage vendor of the node.

Job Name (Type)

The job and job type associated with the Instant Disk Restore.

Recovery Time

The time at which the restore job completed.

Related concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

NetApp ONTAP Aggregates Report

Report the storage utilization and configuration of your aggregates before reaching your aggregate capacity. An aggregate is a collection of disks logically grouped together that provide storage to the volumes that they contain. Run the Aggregates report to view the total and free space available on your aggregate to help you determine if reallocation is necessary based on the size, the number of volumes and disks, and the percentage of space available on your aggregates.

Before you begin

Create and run a NetApp ONTAP Storage Inventory job. You can select one or more NetApp ONTAP providers in a single job definition for cataloging. See [“Creating a NetApp ONTAP Storage Inventory job definition” on page 103](#).

Use the Aggregates report to answer questions such as:

- What is the total size of the aggregates?
- What is the amount of free and used space on the aggregate?
- How many disks make up the aggregate?

Parameters

Use the following parameters to customize your report:

- Storage Array
- Detail View Filter

The default report parameters report on all aggregates with more than 80% space used.

Quick View

The Quick View section displays a pie chart of used and free space on your aggregates. Use the NetApp ONTAP Storage parameter to display aggregates on all storage systems or a specific storage system.

Note: The Quick View section is only modified through the NetApp ONTAP Storage parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the Aggregates report:

Storage Array

The physical server where your files are stored.

Aggregate Count

The number of aggregates on the node.

Volume Count

The number of volumes that make up the aggregate.

Disk Count

The number of disks that make up the aggregate.

Total

The total size of the aggregate.

Available

The amount of free space available in the aggregate.

% Used

The percentage of used storage space on the aggregate.

Detail View

The following fields and corresponding data display in the Detail View section of the Aggregates report:

Aggregate

The name of the aggregate

Storage Array

The physical server where your files are stored.

Location

The node where the volume resides.

Total

The total size of the aggregate.

Available

The amount of free space available in the aggregate.

Volume Count

The number of volumes that make up the aggregate.

Disk Count

The number of disks that make up the aggregate.

Status

The status of the aggregate. An aggregate can be online, offline for maintenance, or reserved for snapshot storage.

% Used/Free

A status bar that displays the used space on the aggregate.

Related concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database

condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

NetApp ONTAP LUNs Report

Review the storage utilization of your NetApp ONTAP LUNs. Run the NetApp ONTAP LUNs report to view your LUNs, total and available space, and online status.

Before you begin

Create and run a NetApp ONTAP Storage Inventory job. You can select one or more NetApp ONTAP providers in a single job definition for cataloging. See [“Creating a NetApp ONTAP Storage Inventory job definition” on page 103](#).

Use the NetApp ONTAP LUNs report to answer questions such as:

- How many LUNs are associated with a storage system?
- What is the total and allocated capacity of a LUN?
- What LUNs are not being used, so I can reclaim this space?

Parameters

Use the following parameters to customize your report:

- Storage Array
- Volume

Multiple selections are supported.

- Detail View Filter

The default report parameters report on all LUNs with more than 80% space used.

Quick View

The Quick View section displays a pie chart of used and free space on your LUNs. Use the NetApp ONTAP Storage parameter to display LUNs on all storage systems or a specific storage system.

Note: The Quick View section is only modified through the NetApp ONTAP Storage parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the NetApp ONTAP LUNs report:

Storage Array

The physical server where your files are stored.

LUNs Count

The number of LUNs on the node.

Total Size

The total size of the aggregate.

Available Size

The amount of free space available in the aggregate.

% Used

The percentage of used storage space on the aggregate.

Detail View

The following fields and corresponding data display in the Detail View section of the NetApp ONTAP LUNs report:

LUN

The name of the LUN.

Storage Array

The physical server where your files are stored.

Volume

The volume associated with the LUN.

Location

The node where the volume resides.

Total Size

The total size of the LUN.

Available Size

The amount of free space available on the LUN.

Status

The status of the LUN. A LUN can be online or offline for maintenance.

Thin Provisioned

The disk format of the LUN, either thick or thin provisioned.

% Used/Free

A status bar that displays the used space on the LUN.

Related concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

NetApp ONTAP Orphaned LUNs Report

Review the NetApp ONTAP Orphaned LUNs Report to view LUNs that have no initiator group mapping or LUNs that belong to volumes that are offline.

Before you begin

Create and run a NetApp ONTAP Storage Inventory job. You can select one or more NetApp ONTAP providers in a single job definition for cataloging. See [“Creating a NetApp ONTAP Storage Inventory job definition” on page 103](#).

Use the NetApp ONTAP Orphaned LUNs report to answer questions such as:

- How much space on an object is consumed by orphaned LUNs?
- Is thin provisioning enabled on a specific LUN?

Parameters

Use the following parameter to customize your report:

- Storage Array
- Volume

Multiple selections are supported.

The default report parameters report on all NetApp ONTAP storage volumes.

Quick View

The Quick View section displays a pie chart of space consumed by orphaned LUNs. Use the NetApp ONTAP Storage parameter to display orphaned LUNs on all storage systems or a specific storage system.

Note: The Quick View section is only modified through the NetApp ONTAP Storage parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the NetApp ONTAP Orphaned LUNs report:

Storage Array

The physical server where your files are stored.

Orphaned LUNs

The number of orphaned LUNs on the node.

Total Size (Volumes)

The total size of the volume on the node.

Total Size (Orphaned LUNs)

The total space on the volume occupied by orphaned LUNs.

% Used (Orphaned LUNs)

The percentage of used storage space on the node.

Detail View

The following fields and corresponding data display in the Detail View section of the NetApp ONTAP Orphaned LUNs report:

LUN

The name of the LUN.

Storage Array

The physical server where your files are stored.

Volume

The volume associated with the LUN.

Location

The node where the volume resides.

Thin Provisioned

The disk format of the LUN, either thick or thin provisioned.

Total Size

The total size of the LUN.

Related concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

NetApp ONTAP Quotas Report

Review the quotas of your NetApp ONTAP storage devices to determine which users or groups are approaching or exceeding their quota limits.

Before you begin

Create and run a NetApp ONTAP Storage Inventory job. You can select one or more NetApp ONTAP providers in a single job definition for cataloging. See [“Creating a NetApp ONTAP Storage Inventory job definition” on page 103](#).

Parameters

Use the following parameter to customize your report:

- Storage Array
- Volume

Multiple selections are supported.

The default report parameters report on all NetApp ONTAP storage volumes.

Quick View

The Quick View section displays a pie chart of space consumed by orphaned LUNs. Use the NetApp ONTAP Storage parameter to display orphaned LUNs on all storage systems or a specific storage system.

Note: The Quick View section is only modified through the NetApp ONTAP Storage parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the NetApp ONTAP Orphaned LUNs report:

Storage Array

The physical server where your files are stored.

Orphaned LUNs

The number of orphaned LUNs on the node.

Total Size (Volumes)

The total size of the volume on the node.

Total Size (Orphaned LUNs)

The total space on the volume occupied by orphaned LUNs.

% Used (Orphaned LUNs)

The percentage of used storage space on the node.

Detail View

The following fields and corresponding data display in the Detail View section of the NetApp ONTAP Orphaned LUNs report:

LUN

The name of the LUN.

Storage Array

The physical server where your files are stored.

Volume

The volume associated with the LUN.

Location

The node where the volume resides.

Thin Provisioned

The disk format of the LUN, either thick or thin provisioned.

Total Size

The total size of the LUN.

Related concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

NetApp ONTAP Snapshots Report

Review the storage utilization of your NetApp ONTAP Snapshots. Run the NetApp ONTAP Snapshots report to view the largest snapshots on your storage systems and the amount of space devoted to snapshot storage.

Before you begin

Create and run a NetApp ONTAP Storage Inventory job. You can select one or more NetApp ONTAP providers in a single job definition for cataloging. See [“Creating a NetApp ONTAP Storage Inventory job definition”](#) on page 103.

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 parameter to customize your report:

- Storage Array
 - Volume
- Multiple selections are supported.
- Number of Largest Snapshots to View

The default report parameters report on the hundred largest snapshots on all volumes.

Quick View

The Quick View section displays a pie chart of the size on your volumes consumed by snapshots. Use the NetApp ONTAP Storage parameter to display snapshots on all storage systems or a specific storage system.

Note: The Quick View section is only modified through the NetApp ONTAP Storage parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the NetApp ONTAP Snapshots report:

Storage Array

The physical server where your files are stored.

Snapshot Count

The number of snapshots on the node.

Total Volume Size

The total size of the volume on which the snapshots are stored.

Total Snapshot Size

The total combined size of all snapshots on the node.

% Used By Snapshot

The percentage of space on the volume used for snapshot storage.

Detail View

The following fields and corresponding data display in the Detail View section of the NetApp ONTAP Snapshots report:

Snapshot

The name of the snapshot.

Storage Array

The physical server where your files are stored.

Volume

The volume on which the snapshot is stored.

Location

The node where the volume resides.

Snapshot Creation Time

The snapshot creation date and time.

Volume Size

The total size of the volume on which the snapshot is stored.

Snapshot Size

The total size of the snapshot.

Total %

The percentage of space on the volume used by the snapshot.

Related concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

NetApp ONTAP Volumes Report

Manage your storage needs and review your volume storage capacity. Run the Volumes report to view the total capacity of your volumes, the total free space, and the available percentage.

Before you begin

Create and run a NetApp ONTAP Storage Inventory job. You can select one or more NetApp ONTAP providers in a single job definition for cataloging. See [“Creating a NetApp ONTAP Storage Inventory job definition” on page 103](#).

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 parameter to customize your report:

- Storage Array
Multiple selections are supported.
- Detail View Filter

The default report parameters report on all NetApp ONTAP storage systems with more than 80% space used.

Quick View

The Quick View section displays a pie chart of used and free space on your volumes. Use the NetApp ONTAP Storage parameter to display volumes on all storage systems or a specific storage system.

Note: The Quick View section is only modified through the NetApp ONTAP Storage parameter.

Summary View

The following fields and corresponding data display in the Summary View section of the NetApp ONTAP Volumes report:

Storage Array

The number of nodes included in the report, based on your parameters.

Volume Count

The number of cataloged volumes included in the report, based on your parameters.

Total

The total space of the volumes included in the report.

Available

The space available on the volumes included in the report.

Reserved

The space reserved for snapshot storage on the volume.

% Used

The percentage of used storage space on the volumes included in the report.

Detail View

The following fields and corresponding data display in the Detail View section of the NetApp ONTAP Volumes report:

Volume

The name of the volume on the node.

Storage Array

The physical server where your files are stored.

Aggregate

The name of the associated aggregate

Location

The node where the volume resides.

Total

The total space on the volume.

Available

The space available on the volume.

Reserved

The space reserved for snapshot storage on the volume.

Status

The online status of the volume.

% Used / Free

The percentage of space used and a status bar that displays the used and free space on the volume. Note that space reserved for snapshot storage is included in this percentage.

Related concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

Pure Storage FlashArray Volumes Report

Report the storage utilization of your Pure Storage FlashArray volumes. Run the Pure Storage FlashArray Volumes report to review the total capacity of your volumes, the total free space, and the percentage available to ascertain your Pure Storage FlashArray volume storage utilization.

Before you begin

Create and run a Pure Storage FlashArray Inventory job. You can select one or more Pure Storage providers in a single job definition for cataloging. See [“Creating a Pure Storage FlashArray Inventory job definition” on page 107](#).

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 Storage 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 concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

Storage Capacity Report

Report the storage capacity of your IBM Storage Virtualize pools, Pure Storage FlashArray volumes, and NetApp ONTAP aggregates. Run the Storage Capacity report to view the total capacity of your volumes, the total free space, and the storage vendor.

Before you begin

- Register storage providers, then create and run Inventory jobs. See [“Register a provider” on page 4](#), and [“Jobs overview” on page 91](#).
- Storage controllers will only be displayed if they have at least one resource successfully backed up to it.

Parameters

Use the following parameters to customize your report:

- Storage Array
- Show Managed Capacity Details
 - Enable to display a detailed view of the managed capacity of storage volumes within a storage array

Summary View

The following fields and corresponding data display in the Summary View section of the Storage Capacity report:

Storage Array

The name of the storage array.

Storage Vendor

The name of the storage vendor of the associated storage provider.

Usable Capacity

The total storage capacity that is available to a storage provider.

Managed Capacity

The space used by IBM Storage Copy Data Management backup jobs on the volumes of an array.

Detail View - Managed Capacity

The Detail View - Managed Capacity section displays if the Show Managed Capacity Details parameter is enabled. The Detail View displays the managed capacity of individual storage volumes within a storage array. The Replication field displays whether a replication relationship is associated with the volume. When calculating the Managed Capacity, the replication relationships are accounted for. If a replication relationship is available, the used size is doubled and displays in the Managed Capacity field.

Related concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

VM and storage mapping report

Review the mappings of virtual machines down to the physical storage from which the datastore is created.

Before you begin

Create and run a VMware Inventory job. You can select one or more VMware providers in a single job definition for cataloging. See [“Creating a VMware Inventory job definition” on page 109](#).

Parameters

Use the following parameters to customize your report:

- vCenter
- ESX Host
- NetApp ONTAP Storage
- Site

Detail View

The following fields and corresponding data display in the Detail View section of the VM Storage and Mapping report:

Datastore

The name of the virtual machine, host, and site and the name of the datastore that is used.

Disk

The disk on which the virtual machine is stored.

Path

The path to the virtual machine disk image file.

LUN

The corresponding logical unit number.

Source NetApp Node : Volume Used/Total

The corresponding NetApp ONTAP Node, volume, and the amount of used and total space.

Replication Destination

If a replication destination exists, it is listed here.

Related concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

VMware Datastores Report

Review the storage utilization of your VMware datastores, including the total free space, provisioned space, and capacities. Run the VMware Datastores report to view your datastores, the number of virtual machines on the datastores, and the percentage of space available.

Before you begin

Create and run a VMware Inventory job. You can select one or more VMware providers in a single job definition for cataloging. See [“Creating a VMware Inventory job definition” on page 109](#).

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 concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

VMware LUNs Report

Review the storage utilization of your VMware LUNs, including the amount of data transferred through the available transport types. Run the VMware LUNs report to view your LUNs, associated datastores, total and allocated capacities, and storage vendors.

Before you begin

Create and run a VMware Inventory job. You can select one or more VMware providers in a single job definition for cataloging. See [“Creating a VMware Inventory job definition” on page 109](#).

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)

Datastore(s)

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 concepts**[“Storage utilization reports” on page 278](#)**

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks**[“Running a report” on page 226](#)**

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.

[“Creating a customized report” on page 227](#)

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.

VMware Orphaned Datastores Report

Review the datastores that do not have any virtual machines assigned to them, or if virtual machines are assigned to the datastores, view the virtual machines that are in an inaccessible state.

Before you begin

Create and run a VMware Inventory job. You can select one or more VMware providers in a single job definition for cataloging. See [“Creating a VMware Inventory job definition” on page 109](#).

Parameters

Use the following parameters 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

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)

Datastore(s)

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 concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

VMware Orphaned LUNs Report

Review VMware orphaned LUNs. These are the LUNs that are neither used as datastores nor RDMs.

Before you begin

Create and run a VMware Inventory job. You can select one or more VMware providers in a single job definition for cataloging. See [“Creating a VMware Inventory job definition” on page 109](#).

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 concepts**[“Storage utilization reports” on page 278](#)**

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks**[“Running a report” on page 226](#)**

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.

[“Creating a customized report” on page 227](#)

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.

VMware VM Snapshot Sprawl Report

The VMware VM Snapshot Sprawl report displays the age and number of snapshots used to protect your virtual machines through NetApp ONTAP Snapshot software..

Before you begin

Create and run a VMware Inventory job. You can select one or more VMware providers in a single job definition for cataloging. See [“Creating a VMware Inventory job definition” on page 109](#).

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 concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

VMware VM Sprawl Report

Review the status of your virtual machines, including virtual machines that are powered off, powered on, or suspended. Run the VMware VM Sprawl report to view unused virtual machines, the date and time they were powered off, and virtual machine templates.

Before you begin

Create and run a VMware Inventory job. You can select one or more VMware providers in a single job definition for cataloging. See [“Creating a VMware Inventory job definition” on page 109](#).

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 concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

VMware VM Storage Report

Review your virtual machines and associated datastores through the VMware VM Storage report. View datastores, resource pools, and the provisioned space of the datastore.

Before you begin

Create and run a VMware Inventory job. You can select one or more VMware providers in a single job definition for cataloging. See [“Creating a VMware Inventory job definition” on page 109](#).

Parameters

Use the following parameters to customize your report:

- vCenter
- ESX Host

Detail View

The following fields and corresponding data display in the Detail View section of the VMware VM Storage report:

VM

The name of the virtual machine, along with the associated datastore.

ESX Host (vCenter)

The host node where the virtual machine resides

Resource Pool

The name of the associated resource pool.

Provisioned Space

The amount of space on the datastore allocated for virtual disk files.

Related concepts

[“Storage utilization reports” on page 278](#)

Storage utilization is a measure of how well your available data storage space is used. The Storage Utilization Reports help you review your storage needs and examine your storage capacity. View the total and free space available as well as the total capacity of your volumes and aggregates. Reports are based on the data collected by the most recently run job.

[“Report Overview” on page 225](#)

IBM Storage Copy Data Management provides a number of predefined reports, which you can tailor to meet your specific reporting requirements. Reports are based on the data collected by the most recently run Inventory job, and you can generate reports after all cataloging jobs and subsequent database condense jobs complete. Click the Reports tab to display the Report Browser. You can run reports with predefined default parameters or run and save customized reports driven by custom parameters.

Related tasks

[“Running a report” on page 226](#)

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.

[“Creating a customized report” on page 227](#)

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.

Chapter 18. Maintenance

The topics in the following section cover maintenance information including updating the IBM Storage Copy Data Management appliance, logging on to the virtual appliance, and collecting logs for troubleshooting.

Maintenance overview

In most cases, IBM Storage Copy Data Management is installed on a virtual appliance. The virtual appliance contains the application and the Inventory.

System Administrators can perform maintenance tasks on the IBM Storage Copy Data Management application. Note that a System Administrator is usually a senior-level user who designed or implemented the vSphere and ESXi infrastructure, or a user with an understanding of IBM Storage Copy Data Management, VMware, and Linux command-line usage. Maintenance tasks are performed in vSphere Client, through the IBM Storage Copy Data Management command-line, or through a web-based management console.

Maintenance tasks include collecting logs, updating the application, and reviewing the configuration of the virtual appliance.

Related concepts

[“Collecting logs for troubleshooting” on page 317](#)

For troubleshooting the IBM Storage Copy Data Management application, IBM Storage Copy Data Management can generate an archive of logs containing various files.

[“Horizontal scale out” on page 331](#)

In some environments, it may make sense to scale out the mongod processes that are part of IBM Storage Copy Data Management. Scale out refers to changing the standard deployment, where the two mongod processes are collocated with IBM Storage Copy Data Management appliance, to a model where each of the MongoDB database instances run in its own virtual machine (VM).

Related tasks

[“Log on to the virtual appliance” on page 316](#)

Log on to the IBM Storage Copy Data Management virtual appliance through vSphere Client to access the command prompt. Available options include collecting logs for troubleshooting.

[“Setting the time zone” on page 316](#)

Access the IBM Storage Copy Data Management Administrative Console to set the time zone of the IBM Storage Copy Data Management appliance.

[“Modifying job log options” on page 320](#)

The settings for job logs can be edited for the IBM Storage Copy Data Management appliance.

[“Managing the administrative console” on page 322](#)

Log on to the Administrative Console to review the configuration of the IBM Storage Copy Data Management virtual appliance. Available information includes general system settings, network, proxy settings, and available updates.

[“Modifying the network settings” on page 329](#)

Access the IBM Storage Copy Data Management appliance using SSH to make adjustments to the network.

[“Uploading an SSL Certificate” on page 330](#)

To establish secure connections in IBM Storage Copy Data Management, you must upload an SSL certificate through the web-based management console of the virtual machine where IBM Storage Copy Data Management is deployed.

[“Restoring a snapshot from a FlexGroup volume to another FlexGroup volume” on page 334](#)

Use this process to manually restore a snapshot from a FlexGroup volume to another FlexGroup volume using the OnTap Command Line Interface (CLI).

Related reference

[“LDAP username syntax” on page 341](#)

LDAP form is used when setting up LDAP authentication in IBM Storage Copy Data Management.

Log on to the virtual appliance

Log on to the IBM Storage Copy Data Management virtual appliance through vSphere Client to access the command prompt. Available options include collecting logs for troubleshooting.

Procedure

To access the virtual appliance command prompt, complete the following steps:

1. In vSphere Client, select the virtual machine where IBM Storage 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`.
4. To log off, enter **exit**.

Related concepts

[“Collecting logs for troubleshooting” on page 317](#)

For troubleshooting the IBM Storage Copy Data Management application, IBM Storage Copy Data Management can generate an archive of logs containing various files.

Related tasks

[“Managing the administrative console” on page 322](#)

Log on to the Administrative Console to review the configuration of the IBM Storage Copy Data Management virtual appliance. Available information includes general system settings, network, proxy settings, and available updates.

Setting the time zone

Access the IBM Storage Copy Data Management Administrative Console to set the time zone of the IBM Storage Copy Data Management appliance.

Procedure

To set the time zone, complete the following steps:

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

```
https://<HOSTNAME>:8090/
```

- where `<HOSTNAME>` is the IP address of the virtual machine where the application is deployed.
2. In the login window, select **System** from the **Authentication Type** drop-down menu. Enter your password to access the Administrative Console. The default password is `ecxadLG235`.
3. Click **Perform System Actions**.
4. In the Change Time zone section, select your time zone. A message stating the operation was successful displays. All IBM Storage Copy Data Management logs and schedules will reflect the selected time zone. The selected time zone will also display on the IBM Storage Copy Data Management appliance when logged in as a root user.
5. To view the current time zone, select **Product Information** from the main Administrative Console page. The displayed time updates every five seconds.

If instances of IBM Storage Copy Data Management reside in time zones that differ from the IBM Storage Copy Data Management appliance, the instances can be synced to the time zone of the IBM Storage Copy

Data Management appliance instead of the time zone of the instance. In some cases, a user may wish to schedule IBM Storage Copy Data Management jobs based on the IBM Storage Copy Data Management appliance's time zone over their local time zone. Perform the following procedure to sync the time zones of the IBM Storage Copy Data Management instances with the IBM Storage Copy Data Management appliance:

6. To sync time zones of IBM Storage Copy Data Management instances with the IBM Storage Copy Data Management appliance, complete the following steps:
 - a) Follow the procedure above to set the time zone of the IBM Storage Copy Data Management appliance.
 - b) Log in to the IBM Storage Copy Data Management appliance as a root user, and edit the following file: `/opt/virgo/repository/ecx-usr/com.syncsort.dp.xsb.api.endeavour.session.properties`.
 - c) In the property file, set `useServerTime` to `true`, then save the file.
 - d) Restart the IBM Storage Copy Data Management appliance through the following commands:

```
service virgo stop
service virgo start
```

Related tasks

[“Log on to the virtual appliance” on page 316](#)

Log on to the IBM Storage Copy Data Management virtual appliance through vSphere Client to access the command prompt. Available options include collecting logs for troubleshooting.

[“Managing the administrative console” on page 322](#)

Log on to the Administrative Console to review the configuration of the IBM Storage Copy Data Management virtual appliance. Available information includes general system settings, network, proxy settings, and available updates.

Collecting logs for troubleshooting

For troubleshooting the IBM Storage Copy Data Management application, IBM Storage 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 Storage Copy Data Management virtual appliance through vSphere Client to download logs using a command prompt. The first approach is simpler and generally sufficient. The second approach produces a more comprehensive set of logs.

Audit logs can also be viewed and downloaded through the **Support** menu. The Audit Log window displays a log of actions performed in IBM Storage Copy Data Management. Information included in the logs include the access time in EPOCH format, IBM Storage Copy Data Management appliance time, IP address of the requester, the username of the requester, the LDAP group to which the user belongs, the operation being performed, and readable text description of the action being performed.

Operations may be one of the following: create, read, update, and delete (CRUD).

Audit Logs can be searched and filtered based on the description, operation, time, or the username of the requester. The filtered log results can be downloaded for future archiving or an external audit.

Collecting audit logs from the support menu

Before you begin

- Contact Technical Support to determine if they need a log collection file for troubleshooting.

Procedure

To collect audit logs from the support menu, complete the following steps:

1. Click the arrow next to the **Support** icon, then click **View Audit Logs**.
2. The Audit Log window displays a log of actions performed in IBM Storage Copy Data Management, along with the user performing the action and a description of the action.
3. To search for the actions of a specific IBM Storage Copy Data Management user, search for the user name in the **Search for users** field.
4. To download the current view of the Audit Log as a .csv file, click **Download**, then select a location to save the file.

What to do next

- Contact Technical Support to inform them that you have created a log collection file for troubleshooting.
- Send the zipped log collection file to Technical Support.
- Manually clean up the archive directory.

Related concepts

[“Monitor a job session” on page 93](#)

You can view the details of a job session that is currently running or one that has finished.

Related tasks

[“Log on to the virtual appliance” on page 316](#)

Log on to the IBM Storage Copy Data Management virtual appliance through vSphere Client to access the command prompt. Available options include collecting logs for troubleshooting.

Collecting logs from the support menu

Before you begin

Contact Technical Support to determine if they need a log collection file for troubleshooting.

Procedure

To collect logs from the support menu, complete the following steps:

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.

What to do next

- Contact Technical Support to inform them that you have created a log collection file for troubleshooting.
- Send the zipped log collection file to Technical Support.
- Manually clean up the archive directory.

Related concepts

[“Monitor a job session” on page 93](#)

You can view the details of a job session that is currently running or one that has finished.

Related tasks

[“Log on to the virtual appliance” on page 316](#)

Log on to the IBM Storage Copy Data Management virtual appliance through vSphere Client to access the command prompt. Available options include collecting logs for troubleshooting.

Collecting the IBM Storage Copy Data Management logs from the virtual appliance

This procedure assumes IBM Storage Copy Data Management deployment was to a VMware appliance host.

Before you begin

- Contact Technical Support to determine if they need a log collection file for troubleshooting.
- Make sure that you have root access to the virtual appliance where IBM Storage Copy Data Management is deployed. See Chapter 3, “[User administration and security management](#),” on page 7.
- 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 [WinSCP](#). If your local computer is Unix based, you can use scp. See [Secure Copy](#).

Procedure

To collect IBM Storage Copy Data Management logs from the virtual appliance, complete the following steps:

1. Log on to the IBM Storage Copy Data Management virtual appliance console as **root** user.
2. Run the below command in the command prompt:

```
cd /opt/ECX/tools/logcollect
```

```
./logcollect
```

3. Enter the **admin** username and password.
4. Run the below command to locate the log collect zip file:

```
cd /opt/ECX/tools/logcollect/logs
```

What to do next

- Contact Technical Support to inform them that you have created a log collection file for troubleshooting.
- copy the zip file to your local computer. If your local computer is Windows based, you can use WinSCP. See [WinSCP](#). If your local computer is Unix based, you can use scp. See [Secure Copy](#).
- Send the zipped log collection file to Technical Support.
- Manually clean up the archive directory.

Related concepts

[“Monitor a job session” on page 93](#)

You can view the details of a job session that is currently running or one that has finished.

Related tasks

[“Log on to the virtual appliance” on page 316](#)

Log on to the IBM Storage Copy Data Management virtual appliance through vSphere Client to access the command prompt. Available options include collecting logs for troubleshooting.

Modifying job log options

The settings for job logs can be edited for the IBM Storage Copy Data Management appliance.

Before you begin

The settings used to control the job logs are listed below with their default values.

- The `maintenance.joblog.enable` setting is a Boolean value that enables or disables the logging of jobs. The default setting for this is `false`.
- The `maintenance.joblog.offload.enable` setting is a Boolean value that enables or disables the offload option for job logs. The default setting for this is `false`.
- The `maintenance.joblog.maxrunningminutes` defines the maximum running time, in minutes, that a job log purge operation will run during a maintenance job. The default setting for this is 60 minutes.
- `Maintenance.joblog.retentiondays` is the number of days that job logs are retained. The default setting for this is 30 days.
- The `maintenance.joblog.export.target` controls the path to which job logs are exported and will only be used if the `maintenance.joblog.offload.enable` option is true. The default path for this is `/data2/joblogs`.

Procedure

To modify job log options, complete the following steps:

1. Log in to the IBM Storage Copy Data Management appliance as the root user.
2. Stop the `virgo` service on the IBM Storage Copy Data Management appliance.

```
$ service virgo stop
```

3. Using a text editor (`vi`), open the `com.syncsort.dp.xsb.serviceprovider.properties` file which is located at `/opt/virgo/repository/ecx-usr/`.
4. Modifying the file as necessary. The settings are listed with their defaults.

```
maintenance.joblog.enable=true
```

```
maintenance.joblog.offload.enable=false
```

```
maintenance.joblog.maxrunningminutes=60
```

```
maintenance.joblog.retentiondays=30
```

```
maintenance.joblog.export.target=/data2/joblogs
```

5. Save the properties file and exit the text editor.
6. Start the `virgo` service on the IBM Storage Copy Data Management appliance.

```
$ service virgo start
```

Updating global settings

You are able to modify a predefined set of configurable options in IBM Storage Copy Data Management through the Update Global Settings window to adjust the settings for your environment.

Before you begin

Ensure that you are aware of the impact of adjusting these settings will have on your environment before making changes.

Note: Values for the settings defined in this topic will not be retained when you update to IBM Storage Copy Data Management from a previous version. If you manually modified any of the settings, you must re-apply those values through the Update Global Settings window after you complete the upgrade by following the procedure in this topic.

Procedure

To modify global settings from the support menu, complete the following steps:

1. Click the arrow next to the **Support** icon, then click **Update Global Settings**.
2. The Update Global Settings window displays a predefined list of options that can be adjusted in your IBM Storage Copy Data Management environment. The settings are divided into three categories: Protection, Recovery, and Job log purge.

Protection

- Protection Incremental FC Copy rate
Enter an integer value to set the copy rate for incremental FlashCopy protection jobs. The default setting is 100 for new deployments.
- Protection Incremental FC Clean rate
Enter an integer value to set the clean rate for incremental FlashCopy protection jobs. The default setting is 100 for new deployments.

Recovery

- Recovery FC Copy rate
Enter an integer value to set the copy rate for FlashCopy recovery jobs. The default setting is 100 for new deployments.
- Recovery FC Clean rate
Enter an integer value to set the clean rate for FlashCopy recovery jobs. The default setting is 0 for new deployments.

Job log purge

- Job log purge
Select this to enable the purging of job logs after a set number of days. When enabled, the value defined in **Job log purge retention** will be used to determine how long to retain job logs before being purged. This is not enabled by default for new deployments.
- Job log purge retention
Enter an integer value to set the number of days to retain job logs. The default setting is 30 days for new deployments. This setting requires that **Job log purge** be enabled.
- Job log purge offload
Select this to enable the offload of job logs to a specific location as defined by the path in the **Job log purge offload target field**. This is not enabled by default for new deployments. This setting requires that **Job log purge** be enabled.
- Job log purge offload target

Enter a path to which job logs will be offloaded. The default setting is `/data2/joblogs` for new deployments. This setting requires that both **Jog log purge** and **Jog log purge offload** be enabled and is required when **Job log purge offload** is enabled.

IBM Storage Copy Data Management Server IP Address

Select the IP address that is assigned to your IBM Storage Copy Data Management appliance. The primary IBM Storage Copy Data Management appliance IP address is the default selectable option in the drop-down menu. If more than one network interface is attached to the IBM Storage Copy Data Management appliance, the IP address assigned to each interface will appear. For example, this can be used to control with which network segment IBM Storage Copy Data Management will interact with if you have an IBM Storage Copy Data Management appliance deployment with an interface mapped to separate network segments. This may also be used to select the IP address if multiple interfaces are available to provide redundancy.

IBM Storage Copy Data Management Server NAT Address

If network address translation (NAT) is used in your environment, enter the NAT addresses or the FQDNs of the server. You must restart IBM Storage Copy Data Management after making changes to this field.

3. Update one or more of the global settings as appropriate to your environment.
4. Click **OK** to save the changes.

Related tasks

[“Managing the administrative console” on page 322](#)

Log on to the Administrative Console to review the configuration of the IBM Storage Copy Data Management virtual appliance. Available information includes general system settings, network, proxy settings, and available updates.

Managing the administrative console

Log on to the Administrative Console to review the configuration of the IBM Storage Copy Data Management virtual appliance. Available information includes general system settings, network, proxy settings, and available updates.

Procedure

To manage the administrative console, complete the following steps:

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

```
https://HOSTNAME:8090/
```

where *HOSTNAME* is the IP address of the virtual machine where the application is deployed.

2. In the login window, select **System** from the **Authentication Type** drop-down menu. Enter your password to access the Administrative Console. The default password is `ecxadLG235`.
3. Review the available options for the virtual appliance.

Pre-upgrade resource adjustments to avoid out-of-memory issues and optimize performance

When you upgrade to IBM Storage Copy Data Management version 2.2.13.0 or later from a previous version, there is a new minimum memory requirement of 48 GB. The memory that is available for Virgo should be set to one quarter of the total memory available. For example, if you use the minimum memory requirement of 48 GB, then 12 GB should be used for Virgo. Follow the steps in this topic to avoid running into any Java™ out-of-memory (OOM) issues and to optimize performance.

About this task

The minimum memory requirements for IBM Storage Copy Data Management version 2.2.13.0 and later have increased. Any Java out-of-memory (OOM) issues should be addressed by following the subsections in the procedure until the issue is resolved.

Procedure

Sub-procedure

The IBM Storage Copy Data Management heap size of 1024 MB may need to be increased to 12288 MB for virtual machine (VM) copy policy catalog browsing. This can be increased by following these steps:

1. Log in to the IBM Storage Copy Data Management virtual appliance through the virtual console or using Secure Shell (SSH) as the root user.
2. Using a text editor, edit the file `startup.sh` file:

```
# vi /opt/virgo/bin/startup.sh
```

3. Locate the line that begins with `export JAVA_OPTS="-Xms512m -Xmx1024m..."` and change it so that it appears as follows:

```
export JAVA_OPTS="-Xms512m -Xmx12288m -XX:PermSize=256m -XX:MaxPermSize=256m"
```

4. Save the file.
5. Reboot the IBM Storage Copy Data Management virtual appliance.

Sub-procedure

If the steps taken in the first sub-procedure do not resolve OOM issues, proceed with the steps in this sub-procedure. To ensure optimal performance, you can optimize the number of concurrent processes within the IBM Storage Copy Data Management virtual appliance by changing the content of the configuration file and restarting Virgo.

6. Log in to the IBM Storage Copy Data Management virtual appliance through the virtual console or using Secure Shell (SSH) as the root user.
7. Using a text editor, edit the file `com.syncsort.dp.xsb.service.jobmanager.properties` file:

```
# vi /opt/virgo/repository/ecx-usr/com.syncsort.dp.xsb.service.jobmanager.properties
```

8. Edit the settings in the file so that they match the values defined below:

```
jobmanager.taskExecutorPoolSize=500
jobmanager.maintenanceTaskExecutorPoolSize=2
jobmanager.threadsPerNode=100
jobmanager.threadsPerJob=100
jobmanager.skipLimit=100
```

9. Save the file.
10. Stop the Virgo service by issuing the stop command:

```
# service virgo stop
```

11. After the service stops, restart it by issuing the start command:

```
# service virgo start
```

Results

The heap of the Virgo component in IBM Storage Copy Data Management virtual appliance is updated to take advantage of increased memory and the number of concurrent processes is updated to optimize performance.

Update IBM Storage Copy Data Management

This topic describes how to update IBM Storage Copy Data Management.

Note: When updating IBM Storage Copy Data Management from version 2.2.19 or earlier through the Administrative Console, a Clean Process Failure message appears. The message indicates that the update applied is successful, but the update cleanup process failed. This message appears consistently when performing an update to IBM Storage Copy Data Management, regardless of whether the Clean Process failed or was successful.

Updating IBM Storage Copy Data Management appliance from 2.2.18 or later versions to 2.2.23 or later

Before you begin

You can update IBM Storage Copy Data Management to the most recent version. If you are using an older version, when you update IBM Storage Copy Data Management you must update it to 2.2.18.0 or 2.2.19.0 and then update to the latest version. For more information about updating IBM Storage Copy Data Management to 2.2.18.0 or 2.2.19.0, see [“Updating IBM Storage Copy Data Management appliance from 2.2.16 or 2.2.17 version to 2.2.18 or later versions” on page 325](#) in the IBM Storage Copy Data Management User's Guide.

- User must restart the virtual machine before updating to IBM Storage Copy Data Management 2.2.23.0.
- Ensure IBM Storage Copy Data Management is installed and running. See [“Installing IBM Storage Copy Data Management as a virtual appliance” on page 16](#) in the IBM Storage Copy Data Management User's Guide.
- After IBM Storage Copy Data Management updates, you 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 Storage Copy Data Management.

Important: Before updating IBM Storage Copy Data Management, you must change the root account password through SSH in the IBM Storage Copy Data Management appliance. SSH into the appliance with the username root and the default password. After authenticating, you will be prompted to change the password from the default.

Procedure

To update your IBM Storage Copy Data Management appliance, complete the following steps:

1. From a machine with internet access, download the update file from IBM [Fix Central](#).
2. From a supported web browser, access the Administrative Console at the following address:

```
https://<HOSTNAME>:8090/
```

where *<HOSTNAME>* is the IP address of the virtual machine where the application is deployed.

3. In the login window, select **System** from the **Authentication Type** drop-down menu. Enter your password to access the Administrative Console. The default password is `ecxadLG235`.
4. Click **Manage updates**. In the **Don't want to update IBM Storage Copy Data Management to the latest available release?** section, select **Click here**.
5. Click **Choose File**, browse for the update file to upload to the appliance, then click **Upload Update Image**. The update process begins after the update image has been uploaded to the appliance.
6. After the update completes, navigate to the **Perform System Actions** page on the Administrative Console to restart the appliance.
7. HTML content from previous versions of IBM Storage Copy Data Management may be stored in your browser's cache. Clear your browser's cache before logging in to an updated version of IBM Storage Copy Data Management to ensure you are viewing the latest content changes.

Remember:

- Updating the IBM Storage Copy Data Management appliance will require that the virtual machine be restarted for plugins to function properly.
- Ignore the **Clean process Failure** error while updating the IBM Storage Copy Data Management appliance.

What to do next

- 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 tasks

[“Installing IBM Storage Copy Data Management as a virtual appliance” on page 16](#)

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

[“Managing the administrative console” on page 322](#)

Log on to the Administrative Console to review the configuration of the IBM Storage Copy Data Management virtual appliance. Available information includes general system settings, network, proxy settings, and available updates.

Updating IBM Storage Copy Data Management appliance from 2.2.16 or 2.2.17 version to 2.2.18 or later versions

Use this procedure to update IBM Storage Copy Data Management appliance version 2.2.16 or 2.2.17 to 2.2.18 or later version.

Before you begin

To update the IBM Storage Copy Data Management appliance to version 2.2.18 or later, make sure that the current IBM Storage Copy Data Management appliance version is either 2.2.16 or 2.2.17. If not update the IBM Storage Copy Data Management appliance previous version to 2.2.16, refer to [Updating IBM Storage Copy Data Management appliance from previous version to 2.2.16](#).

Important: Before updating IBM Storage Copy Data Management, you must change the root account password through SSH in the IBM Storage Copy Data Management appliance. SSH into the appliance with the username root and the default password. After authenticating, you will be prompted to change the password from the default.

Procedure

Run a script manually to update IBM Storage Copy Data Management appliance version 2.2.16 or 2.2.17 to 2.2.18 or later:

1. Take a vmware snapshot of your IBM Storage Copy Data Management virtual appliance.
2. Login to the IBM Storage Copy Data Management virtual appliance as the root user and download the 2.2.18 or later ISO.
3. Download the 2.2.18 or later ISO to the required folder for example: `/tmp`.
4. Run a script that follows:

```
mkdir -p /mnt/SCDMiso
```

```
mount -o loop /tmp/<scdm ISO> /mnt/SCDMiso
```

```
cd /mnt/SCDMiso
```

```
./upgradeSCDM.sh
```

Note: After updating to 2.2.18 or later using the script, you must reboot the appliance.

Updating IBM Storage Copy Data Management appliance from 2.2.15 or earlier versions to 2.2.16 or later versions

Before you begin

You can update IBM Storage Copy Data Management to the most recent version. When you update IBM Storage Copy Data Management, you must update from no more than two previous versions (n-2) to the current version (n). If you are using an older version, you must update at least to (n-2) version and then update to the current version. When you update to IBM Storage Copy Data Management 2.2.16.0, you must complete a new installation of IBM Storage Copy Data Management and then migrate from the previous version. This requires the catalog to be transferred from the previous version of IBM Storage Copy Data Management to the new deployment of IBM Storage Copy Data Management 2.2.16.0. For more information about the catalog migration process, follow the procedure in [“Transferring the catalog from a previous version of IBM Storage Copy Data Management”](#) on page 326.

- Ensure IBM Storage Copy Data Management is installed and running. See [“Installing IBM Storage Copy Data Management as a virtual appliance”](#) on page 16.
- After IBM Storage 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 Storage Copy Data Management.

Important: Before updating IBM Storage Copy Data Management, you must change the root account password through SSH in the IBM Storage Copy Data Management appliance. SSH into the appliance with the username root and the default password. After authenticating, you will be prompted to change the password from the default.

What to do next

- 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 tasks

[“Installing IBM Storage Copy Data Management as a virtual appliance”](#) on page 16

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

[“Managing the administrative console”](#) on page 322

Log on to the Administrative Console to review the configuration of the IBM Storage Copy Data Management virtual appliance. Available information includes general system settings, network, proxy settings, and available updates.

Transferring the catalog from a previous version of IBM Storage Copy Data Management

The catalog is critical to the operation of IBM Storage Copy Data Management. The catalog must be transferred from previous versions of IBM Storage Copy Data Management to IBM Storage Copy Data

Management 2.2.21 when you upgrade. This procedure is required when you upgrade to IBM Storage Copy Data Management 2.2.21.

About this task

In IBM Storage Copy Data Management 2.2.21, the platform was upgraded from previous releases. There is not a direct upgrade path from previous versions of IBM Storage Copy Data Management to version 2.2.21 as it is now based on Redhat Linux. As a result, you must deploy a new instance of IBM Storage Copy Data Management 2.2.21 and transfer the catalog from the previous version of IBM Storage Copy Data Management and restore it in the newly deployed version.

Note: If you are updating the IBM Storage Copy Data Management version from 2.2.15 or earlier to 2.2.16 or later versions, you must transfer the catalog.

Procedure

1. Using a supported web-browser, log in to your existing IBM Storage Copy Data Management virtual appliance.
2. Prevent any jobs from running by selecting **Hold Schedule** for all backup, restore, inventory and maintenance jobs that have the schedule enabled. For more information on holding jobs, see [“Holding and releasing a job session” on page 92](#) in the IBM Storage Copy Data Management User's Guide.

Important:

- Jobs using scripts do not require any special action when updating with catalog backup/restore from IBM Storage Copy Data Management as they can be edited/saved after the update.
 - If you're updating from IBM Storage Copy Data Management 2.2.13.0 to 2.2.21 for this release, active session control is only an issue that should be canceled or moved to permanent before the update. Updates from IBM Storage Copy Data Management 2.2.14.0/2.2.15.0/2.2.16.0/2.2.17.0/2.2.18.0/2.2.19.0/2.2.20.0 to IBM Storage Copy Data Management 2.2.21 will require no special action for Active Session control.
 - You must cancel or move to permanent all **Resource Active** sessions before you upgrade to IBM Storage Copy Data Management from a previous version. Jobs from previous versions of IBM Storage Copy Data Management will not load former active sessions, because the format of the data in the recovery session objects changes as part of the upgrade process.
3. Log in to the Administrative Console of the existing IBM Storage Copy Data Management virtual appliance.
 4. Click the **Menu > Catalog Manager > Catalog Backup > Backup**. The catalog backup will commence. Wait for the backup operation to complete.
 5. Next, SSH into the existing IBM Storage Copy Data Management virtual appliance. Copy the contents of `/data/backup` to another location separate from the existing IBM Storage Copy Data Management virtual appliance that can be transferred when needed in a later step.
 6. Deploy IBM Storage Copy Data Management 2.2.21. For information on deploying IBM Storage Copy Data Management, refer to [“Installing IBM Storage Copy Data Management as a virtual appliance” on page 16](#) in the IBM Storage Copy Data Management User's Guide.
 7. Change the default passwords on the newly deployed IBM Storage Copy Data Management virtual appliance. You must also change the password for the `root` user from the virtual appliance's command line via SSH. For more information, see [“Access and default credentials” on page 12](#) in the IBM Storage Copy Data Management User's Guide.

Note: You must change the **SYSTEM** account password through SSH before you can access the Administrative Console. SSH into the appliance with the username `administrator` and the default password. After authenticating, you will be prompted to change the password from the default. This account can now be used to log into the Administrative Console.

8. Using SSH, log in to the UI of the newly deployed IBM Storage Copy Data Management virtual appliance as the `root` user. Create a backup directory as a sub-directory in `/data`:

```
mkdir /data/backup
```

9. Copy the data contained in the /data/backup directory that you copied in Step 5 into the backup directory created in the step 8.
10. Log in to the Administrative Console of the newly deployed IBM Storage Copy Data Management virtual appliance. Click **Menu** and then click **Catalog Manager**.
11. Click the **Menu > Catalog Manager > Catalog Restore > Restore**. Wait for the restore operation to complete.
12. Log in to the newly deployed IBM Storage Copy Data Management UI. Release the schedule for all jobs by selecting **Release Schedule** for all jobs. For more information on releasing jobs, refer to [“Holding and releasing a job session” on page 92](#).
13. Re-import the source IBM Storage Copy Data Management license into the target IBM Storage Copy Data Management unless the source IBM Storage Copy Data Management and target IBM Storage Copy Data Management both have permanent licenses. If the source IBM Storage Copy Data Management license is permanent and the target IBM Storage Copy Data Management license is trial, the source IBM Storage Copy Data Management license should be imported into the target IBM Storage Copy Data Management.
14. Power off the older IBM Storage Copy Data Management virtual appliance. It may be beneficial to keep this appliance in the powered off state for a period of time in the case that it is needed for reference.

Results

The catalog from the previous version of IBM Storage Copy Data Management is migrated to the newly deployed version of IBM Storage Copy Data Management.

Backup and Restore the Catalog

The Catalog Manager provides IBM Storage Copy Data Management administrators with the ability to backup and restore the IBM Storage Copy Data Management Catalog.

Procedure

To manage your IBM Storage Copy Data Management catalog through the Catalog Manager, complete the following steps:

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

```
https://<HOSTNAME>:8090/
```

where *<HOSTNAME>* is the IP address of the virtual machine where the application is deployed.

2. In the login window, select **System** from the **Authentication Type** drop-down menu. Enter your password to access the Administrative Console. The default password is `ecxadLG235`.
3. Click **Menu**, then select **Catalog Manager**.
4. Select Backup Catalog or Restore Catalog.

Backup Catalog: In the Directory field, enter the backup destination on the IBM Storage Copy Data Management host. Ensure the destination volume exists and that there is enough room on the destination volume for the Catalog backup. Click Backup to begin the Catalog backup.

Backup Catalog considerations: IBM Storage Copy Data Management will be stopped while the Catalog is being backed up. The IBM Storage Copy Data Management user interface will not be accessible, and all running jobs will be aborted.

Restore Catalog: In the Directory field, enter the source of the Catalog restore on the IBM Storage Copy Data Management host. Ensure the source volume exists. Click Restore to begin the Catalog restore.

Restore Catalog considerations: IBM Storage Copy Data Management will be stopped while the Catalog is being restore. The IBM Storage Copy Data Management user interface will not be accessible, and all running jobs will be aborted. All IBM Storage Copy Data Management snapshots created after the Catalog backup was run will be lost.

Related tasks

[“Managing the administrative console” on page 322](#)

Log on to the Administrative Console to review the configuration of the IBM Storage Copy Data Management virtual appliance. Available information includes general system settings, network, proxy settings, and available updates.

Modifying the network settings

Access the IBM Storage Copy Data Management appliance using SSH to make adjustments to the network.

Procedure

To modify the network settings, complete the following steps:

Users can configure the network post deployment should it be required. Follow this procedure to utilize the NetworkManager Text User Interface (nmtui) tool. When connecting to IBM Storage Copy Data Management to use this tool, the `root` account or an account with root privileges will be required.

1. SSH from a terminal to the IBM Storage Copy Data Management appliance:

```
ssh root@appliancehostname
```

2. Enter password for the `root` account or the password associated with the username. For more information, see [“Access and default credentials” on page 12](#).
3. After you have successfully connected, enter the following to launch the tool:

```
nmtui
```

4. Edit the network configuration:

Edit a connection

- a) Select `ens160` from the list and select **<Edit...>**.
- b) Here you can edit the IPv4 configuration. Set the connection **Automatic** to use DHCP or **Manual** to use static IP addressing. You can add **Addresses**, **DNS servers**, and **Search domains**. You may also specify the **Gateway** IP address.
- c) After completing changes, select **<OK>**.
- d) Click **<Back>**.

Set system hostname

- a) Enter a valid hostname in the **Hostname** field.
- b) Click **<OK>**.
- c) On the Set hostname to screen select **OK** to verify the change.

5. Select **<Quit>** and then select **<OK>**.

6. To update VADP proxy server settings after an IBM Storage Copy Data Management IP or hostname change, complete the following steps:

Changes made to the IP or hostname of the IBM Storage Copy Data Management appliance will result in a loss of communication with associated VADP proxy servers. Please follow the steps below on each associated VADP proxy server to re-establish communication with the IBM Storage Copy Data Management appliance.

SSH to each VADP proxy and enter the following commands:

- a) Enter the new IBM Storage Copy Data Management appliance IP or hostname to the ECX_HOST variable in /opt/ECX/bin/escvadvp and save:

```
vi /opt/ECX/bin/escvadvp
ECX_HOST=<new ECX IP or Hostname>
```

- b) Enter the new IBM Storage Copy Data Management appliance IP or hostname to the ECX_HOST variable in /etc/rc.d/init.d/escvadvp and save:

```
vi /etc/rc.d/init.d/escvadvp
ECX_HOST=<new ECX IP or Hostname>
```

- c) Restart the ecxvadvp service:

```
service ecxvadvp restart
```

Any associated VADP proxy servers should be now able to communicate with the IBM Storage Copy Data Management appliance using the updated IP address or hostname.

Related tasks

[“Log on to the virtual appliance” on page 316](#)

Log on to the IBM Storage Copy Data Management virtual appliance through vSphere Client to access the command prompt. Available options include collecting logs for troubleshooting.

[“Managing the administrative console” on page 322](#)

Log on to the Administrative Console to review the configuration of the IBM Storage Copy Data Management virtual appliance. Available information includes general system settings, network, proxy settings, and available updates.

Uploading an SSL Certificate

To establish secure connections in IBM Storage Copy Data Management, you must upload an SSL certificate through the web-based management console of the virtual machine where IBM Storage Copy Data Management is deployed.

Before you begin

- If you are uploading an LDAP SSL certificate, ensure an LDAP server is running and reachable by IBM Storage Copy Data Management.
- If you are uploading an LDAP SSL certificate, review LDAP syntax. See [“LDAP username syntax” on page 341](#).
- To import the HTTP certificate, enter the key store password in the password field. The default key store password is ecx-beta.

Procedure

To upload a certificate, complete the following steps:

1. Contact your network administrator for the name of the certificate to export.
2. From a supported browser, export the certificate to your computer. Make note of the location of the certificate on your computer. The process of exporting certificates varies based on your browser. See Related Topics.
3. From a supported browser, enter the following URL:

```
https://<HOSTNAME>:8090/
```

where <HOSTNAME> is the IP address of the virtual machine where the application is deployed.

4. In the login window, select **System** from the **Authentication Type** drop-down menu. Enter your password to access the Administrative Console. The default password is ecxadLG235.
5. Click **Manage your certificates**.

6. Select the Certificate Type, browse for the certificate file on your computer.
7. Click the password field and enter the key store password to import the certificate. The default key store password is ecx-beta.
8. Click **Upload**.
9. Reboot the virtual machine where the application is deployed.

Related concepts

[“User administration and security management” on page 7](#)

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

Related tasks

[“Register a provider” on page 4](#)

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

Related reference

[“LDAP username syntax” on page 341](#)

LDAP form is used when setting up LDAP authentication in IBM Storage Copy Data Management.

Related information

[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](#)

Horizontal scale out

In some environments, it may make sense to scale out the mongod processes that are part of IBM Storage Copy Data Management. Scale out refers to changing the standard deployment, where the two mongod processes are collocated with IBM Storage Copy Data Management appliance, to a model where each of the MongoDB database instances run in its own virtual machine (VM).

The general steps to accomplish this are as follows:

- Deploy a MongoDB virtual machine (VM) from the .ova file.
- Copy the appropriate VMDKs from IBM Storage Copy Data Management to the MongoDB VM.
- Configure the MongoDB VM.

Note: Tasks Deploy a MongoDB virtual machine (VM) from the .ova file, Copy the appropriate VMDKs from IBM Storage Copy Data Management to the MongoDB VM, and Configure the MongoDB VM will need to be performed twice, once for each of the two MongoDB databases. Proceed to the task, Configure IBM Storage Copy Data Management to use the MongoDB VMs rather than the collocated MongoDB processes only after these tasks have been done for each database.

- Configure IBM Storage Copy Data Management so that it will use the MongoDB VMs as opposed to the collocated mongod processes running the MongoDB databases.
- Bind directories and edit the fstab file to ensure proper function of log collection.

Before you begin

- Obtain the VMware .ova that deploys a MongoDB VM from Catalogthe **<xxx>**.
- Locate the PowerShell script named copy_mongodb_vmdk.ps1 that copies the VMDKs from IBM Storage Copy Data Management that is included in the .ova.

- PowerShell 5.1 or later or PowerCLI 11.2 or later are required to use the `copy_mongodb_vmdk.ps1` script.
- The Node.js script (included in the `.ova`) that manages the MongoDB VM.
- Ensure the account with which you are accessing the vSphere UI has the **Datastore.Browse** privilege.
- Determine two available static IPv4 addresses that can be used for the two VMs that will be deployed.
- If updating from IBM Storage Copy Data Management 2.9 to IBM Storage Copy Data Management 2.10, the **nfs-utils** module needs to be installed. Run the following command from an IBM Storage Copy Data Management appliance terminal window as the **root** user:

```
$ sudo yum --enablerepo=base,updates install -y nfs-utils
```

Deploying a MongoDB virtual machine (VM) from the .ova file

This procedure tells you how to deploy a MongoDB virtual machine from the `.ova` file

Procedure

1. Log in to the vSphere UI
2. Deploy the MongoDB VM `.ova` file. DO NOT power on the VM.

Copying the appropriate VMDKs from IBM Storage Copy Data Management to the MongoDB VM

This procedure describes how to Copy the appropriate VMDKs from IBM Storage Copy Data Management to the MongoDB VM

Before you begin

- Ensure that both IBM Storage Copy Data Management and the MongoDB VM are powered off.
- Have available the IP address of the vCenter server that manages IBM Storage Copy Data Management appliance.
- Determine the names of the IBM Storage Copy Data Management and MongoDB VM. These are the VM names displayed in the vCenter UI.

Procedure

1. Open a PowerShell session.
2. In the PowerShell session, navigate (`cd`), into the directory that contains the PowerShell script, `copy_mongodb_vmdk.ps1`.
3. Enter the command, substituting the vCenter IP for **<vCenter_ip>**, the vCenter user name for **<vCenter_userid>**, the IBM Storage Copy Data Management VM name for **<IBM Storage Copy Data Management_VM_name>**, and the MongoDB VM name for **<MongoDB_VM_name>**:

```
./copy_mongodb_vmdk.ps1 mongo1 <vCenter_ip> <vCenter_userid> <IBM Storage Copy Data Management_VM_name> <MongoDB_VM_name> -whatif:$false
```

4. The script will prompt for the hard disk number that corresponds to the database (ex. `mongo1`). Type the number at the prompt and press **Enter**.
5. When all hard disk numbers have been entered, type the letter **'q'** at the prompt.

Note: Depending on the size of the VMDKs and the speed of the storage network, the script may run for as long as one hour. To observe progress, consult the vSphere UI.

Configuring the MongoDB VM

This procedure describes how to Configure the MongoDB VM

Before you begin

- A deployed, powered off MongoDB VM to which the IBM Storage Copy Data Management VMDK(s) have been copied.
- Have available the IP address of the IBM Storage Copy Data Management virtual appliance.

Procedure

1. Log in to the vSphere UI.
2. Power on the MongoDB VM.
3. After the MongoDB VM is up, use the vSphere UI console facility to open the VM's console.
4. Log in to the VM's shell using a user name of **ecxadmin** and the password **ecxDP758**. If prompted, set a new password.
5. Edit (using vi) the file `config.yml` and set the property: IBM Storage Copy Data Management_ip: <IBM Storage Copy Data Management_ip> where **<IBM Storage Copy Data Management_ip>** is the IP of the IBM Storage Copy Data Management appliance.
6. In the shell, enter the command: `sudo ./mongoman --func=dsc`. The "Starting MongoDB... done" message will appear.
7. Enter the command: `sudo ./mongoman`. Based on the VM being configured (mongo1 or mongo2), the "Database Name: mongo1(up)" or "Database Name: mongo2(up)".

Steps 1 - 3 need to be repeated for the second database, mongo2.

Tip: Remember to adjust the database names from mongo1 to mongo2 for the **<MongoDB_VM_name>** variable with executing the `copy_mongodb_vmdk.ps1` PowerShell script and adjust the hard disk number that corresponds to mongo1 to that of mongo2.

Proceed with Step 4 only be after Steps 1 - 3 have been completed for both mongo1 and mongo2.

Configuring IBM Storage Copy Data Management to use the MongoDB VMs rather than the collocated MongoDB processes

This procedure describes how to Configure IBM Storage Copy Data Management to use the MongoDB VMs rather than the collocated MongoDB processes.

Before you begin

- Steps 1 - 3 have been completed for both mongo1 and mongo2.
- A powered off IBM Storage Copy Data Management appliance that is in its default configuration with collocated MongoDBs.
- The password for the IBM Storage Copy Data Management appliance's root account.
- Have available the IP addresses of the mongo1 and mongo2 VMs.

Procedure

1. Log in to the vSphere UI.
2. Power on the IBM Storage Copy Data Management VM through the vSphere UI.
3. After the IBM Storage Copy Data Management VM is up, use the vSphere UI console facility to open the VM's console.
4. Log in to the IBM Storage Copy Data Management shell as **root** using the password set when you first logged into the IBM Storage Copy Data Management shell.

5. Edit (using vi) the managedb.config file located at

/opt/ECX/tools/scripts/managedb.config and set the properties:

```
MONGO_HOST=<mongo1_ip>
MONGO2_HOST=<mongo2_ip>
MONGO2_PORT=27017
```

where **<mongo1_ip>** is the IP of mongo1 and **<mongo2_ip>** is the IP of mongo2.

6. Restart the IBM Storage Copy Data Management appliance. Please give IBM Storage Copy Data Management several minutes to fully restart.

Binding directories to ensure proper function of log collection

This procedure describes how to bind directories to ensure proper function of log collection.

About this task

On the IBM Storage Copy Data Management appliance, create and bind the directories that are required for log collection to the newly created mongo1 and mongo2 disks.

Procedure

1. Log in to the IBM Storage Copy Data Management appliance.
2. Open a terminal and enter the following commands, where **<mongo1_ip>** is the IP address of mongo1, and **<mongo2_ip>** is the IP address of mongo2:

```
$ sudo mkdir /tmp/mongo1
$ sudo mount -t nfs <MONGO1_IP>:/data/log/mongo /tmp/mongo1
$ sudo mkdir /tmp/mongo2
$ sudo mount -t nfs <MONGO2_IP>:/data/log/mongo /tmp/mongo2
```

3. To ensure persistence after appliance restart, edit fstab to mount the directories automatically. Use vi to edit the fstab file:

```
$ sudo vi /etc/fstab
```

4. Once the fstab file is open, append the following lines:

```
<mongo1_ip>:/data/log/mongo /tmp/mongo1 nfs _netdev 0 0
<mongo2_ip>:/data/log/mongo /tmp/mongo2 nfs _netdev 0 0
```

5. Save changes to the fstab file and exit.

Results

It may take several minutes for IBM Storage Copy Data Management to fully restart. Once IBM Storage Copy Data Management has restarted, log in to the IBM Storage Copy Data Management appliance UI to ensure that the procedure completed successfully.

Restoring a snapshot from a FlexGroup volume to another FlexGroup volume

Use this process to manually restore a snapshot from a FlexGroup volume to another FlexGroup volume using the OnTap Command Line Interface (CLI).

Before you begin

Consider the following requirements before attempting this procedure.

- Perform the process as a cluster administrator at the **admin** privilege level. If you do not have access to the cluster admin account, using a cluster administrative account with similar access permissions will be necessary.

- The `cifs.show_snapshot` option needs to be set to on. This can be verified from the OnTap CLI on the NetApp controller by issuing:

```
options cifs.show_snapshot
```

- If it is not enabled (off), turn it on by issuing the following command:

```
options cifs.show_snapshot on
```

- For specific FlexGroup requirements, please review the NetApp OnTap Documentation. See the information at [ONTAP 9 Documentation Center](#).

About this task

If it has not already done so, mount the CIFS share that contains the snapshots to be restored. The `.snapshot` directory will appear at the root of the share.

Note: When viewing contents of the home directory, the directory may be displayed in one of two ways. Likely, if long file names are supported, the directory will be displayed as `~snapshot`, whereas if the operating system does not support long file names, the directory will be displayed as `~SNAPSHT`.

Procedure

To restore a snapshot from a FlexGroup volume to another FlexGroup volume, complete the following steps:

1. Open the OnTap Command Line Interface (CLI).
2. Obtain the name of the source and destination volumes, and the name of the snapshot to be restored.
3. Deactivate any active quota rules on the destination FlexGroup volume. Active quota rules can be deactivated by issuing the following command from the OnTap CLI:

```
volume quota modify -vserver <vserver> -volume <volume> -state off
```

Note: Information contained between the less-than and greater-than symbols should contain the source and destination servers and flexgroups. Omit the less-than and greater-than symbols when running the command. The colon and dashes should remain.

4. Enter the following command to initiate the restore:

```
snapmirror restore -source-path <source vserver>:<source flexgroup> -destination-path <destination vserver>:<destination flexgroup> -snapshot <snapshot name>
```

5. If there were active quota rules applied to the destination FlexGroup volumes that were deactivated in Step 3, reactivate them. Use the following command to reactivate quote rules:

```
volume quota modify -vserver <vserver> -volume <volume> -state on
```

Chapter 19. Documentation and support

The topics in the following section cover the documentation roadmap, technical support details, and information about the online help system.

Documentation roadmap

Help System

In IBM Storage Copy Data Management, when needed:

- Click the help icon to invoke help specific to the active function.
- Use the Help system's Search and Index features to locate pertinent information, as these features search the entire documentation suite.

User's Guide

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

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

Related concepts

[“About the help system” on page 337](#)

About the help system

Starting Help

In the application, click the help icon to invoke help specific to the active function. For example, click the help icon on the Search tab to view help topics related to searching for objects.

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

Before You Begin, Next Steps, and Related Topics

Prerequisites for procedures are listed in Before You Begin sections in many help topics. For example, you must run a job before you run a report, so a link to the Run a Job procedure is available in the Before You Begin section of the Run a Report procedure. Topics also include Next Steps and Related Topics sections for more information.

Search Help Feature

Use the Search feature in the Help system to locate pertinent information in the entire documentation suite:

- Enter a word in the search field to find all topics that contain that word. For example, schedule.
- Enter multiple words in the search field to find topics related to both words. For example, searching for schedule job returns results for schedule and job as well as "schedule job."
- Enter words separated with AND, +, or & to find topics that contain all of the words separated by the operators. For example, schedule AND job, schedule + job, or schedule & job.

- Enter words separated with OR to find topic that contain any of the words separated by OR. For example, schedule OR job.

Search for words on a help page by using the Find feature in your browser.

Security Management Topics

A security icon precedes a security management help topic. Security management identifies the interfaces that manage the security functions in IBM Storage 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 Storage Copy Data Management to use LDAP, and configuring IBM Storage Copy Data Management to use HTTPS.

Related concepts

[“Documentation roadmap” on page 337](#)

Chapter 20. Reference topics

The topics in the following section cover reference information including virtual machine privileges, search and filter guidelines, and return code references.

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 tasks

[“Search for objects” on page 4](#)

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

[“Viewing object details” on page 218](#)

From the search results, you can view the attributes of a searched object, including its location, type, and the dates associated with its creation and modification. If more than one version of an object exists in the Inventory, the search results displays the latest version and attributes of the object.

[“Viewing NetApp ONTAP file details” on page 219](#)

From the search results, you can view the attributes of a searched file, including its location, size, and the date the file was added to the Inventory. You can also quickly view a file's creation time, last modified time, and the last accessed time from the search results pane.

[“Creating a NetApp ONTAP File Inventory job definition” on page 104](#)

A NetApp ONTAP File Inventory job provides the framework to catalog and collect information about low-level NetApp ONTAP objects.

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 concepts

[“Viewing a provider” on page 40](#)

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.

[“Monitor a job session” on page 93](#)

You can view the details of a job session that is currently running or one that has finished.

[“Searching overview” on page 213](#)

Use IBM Storage 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.

Related tasks

[“Editing a schedule” on page 88](#)

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.

[“Editing a job definition” on page 211](#)

The procedures describe how to edit a job definition.

[“Running a report” on page 226](#)

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.

LDAP username syntax

LDAP form is used when setting up LDAP authentication in IBM Storage 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 concepts

[“User administration and security management” on page 7](#)

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

Related tasks

[“Uploading an SSL Certificate” on page 330](#)

To establish secure connections in IBM Storage Copy Data Management, you must upload an SSL certificate through the web-based management console of the virtual machine where IBM Storage Copy Data Management is deployed.

Return code reference

Return Codes are issued when a script generated by a user-defined script is run. As the script runs, IBM Storage 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 (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

Return Code (128+N)	Return Code Value
19	SIGSTOP
20	SIGTSTP
21	SIGTTIN
22	SIGTTOU
23	SIGURG
24	SIGXCPU
25	SIGXFSZ
26	SIGVTALRM
27	SIGPROF
28	SIGWINCH
29	SIGIO
30	SIGPWR
31	SIGSYS
32	
33	
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
48	SIGRTMIN+14
49	SIGRTMIN+15
50	SIGRTMAX-14
51	SIGRTMAX-13
52	SIGRTMAX-12
53	SIGRTMAX-11

Return Code (128+N)	Return Code Value
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 tasks

[“Creating a script job definition” on page 209](#)

The procedures describe how to create a script job definition.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Red Hat, OpenShift®, Ansible®, and Ceph® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

VMware, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Glossary

A glossary is available with terms and definitions for the IBM Storage Copy Data Management family of products.

See the IBM Storage Copy Data Management [glossary](#).



Product Number: 5737-B34
5641-CD4
5641-CD5
5641-CD6