

IBM Storage Defender Copy Data Management

Installation and User's Guide

2.2.28



Contents

Audience and purpose	10
IBM® Storage Defender Copy Data Management enhancements	11
IBM® Storage Defender Sentinel	11
Limitations and known issues	11
Common limitations and known issues	12
SAP HANA database specific limitations	14
Oracle database specific limitations	15
Microsoft SQL Server specific limitations	15
VMware specific limitations	15
Dell PowerMax Storage specific limitations	15
Dell PowerFlex Storage specific limitations	16
IBM Storage Virtualize for Snapshot specific limitations	16
Policy-based high availability (PBHA) specific limitations	17
Policy-based replication (PBR) specific limitations	18
Sentinel specific limitations	18
IBM® Storage Defender Copy Data Management overview	20
Getting off to a quick start	21
Start IBM® Storage Defender Copy Data Management	21
Register a provider	21
Create an inventory job definition	22
Run a job	22
Search for objects	22
Generate reports	22
Backup	23
Restore	23
User administration and security management	24
Identification and Authentication	24
User Data Security	24
Security Management	24
Management and Operation Functions	25
Encryption	25
Ports	25
Installation and setup	27
Deployment checklist	27
Access and default credentials	28
VMware vSphere™ Privileges	30
Installing IBM® Storage Defender Copy Data Management as a virtual appliance	32
Installing IBM® Storage Defender Sentinel Security Scan	34
Uninstalling the agent from Windows™-based application server	34
Starting IBM® Storage Defender Copy Data Management	35
Dashboard	36
Available Dashboard Widgets	36
Sites and providers	38
Sites and providers overview	38
Adding a site	39
Editing a site	39
Deleting a site	40
Registering a storage provider	40
Registering an Application Server - File System	40
Registering an Application Server - InterSystems Database	41
Registering an Application Server - Oracle	43
Registering an Application Server - SAP HANA	44
Registering an Application Server - SQL	45
Registering an IBM® Storage Virtualize provider	47
Registering an IBM Storage Virtualize for Snapshot provider	48
Registering an LDAP provider	49
Registering a NetApp ONTAP provider	50
Registering multiple NetApp ONTAP providers through the Discover feature	51
Registering a Pure Storage FlashArray provider	52
Registering a Security Scan Server in IBM® Storage Defender Copy Data Management	53

Registering an SMTP provider.....	54
Registering a VMware provider.....	55
Registering multiple VMware providers through the Discover feature.....	56
Registering a Dell PowerMax provider.....	57
Registering a Dell PowerFlex Storage provider.....	58
Viewing a provider.....	60
Viewing a list of providers.....	61
Editing a provider.....	61
Unregistering a provider.....	62
Adding credentials to a virtual machine.....	63
Adding credentials for a single virtual machine.....	63
Adding credentials for multiple virtual machines.....	64
Requiring multifactor authentication.....	65
Multifactor authentication for the ADMIN user.....	65
Creating multifactor authentication user account.....	66
Enabling time-based one-time multifactor authentication for a user.....	66
Setting up a multifactor authentication user account in the browser.....	66
Expiring time-based one-time secret key of a user account.....	67
Disabling multifactor authentication.....	67
Role-based access control.....	68
Role-based access control overview.....	68
Configure resource pools.....	69
Configure roles.....	71
Configure accounts.....	72
VMware admin role-based access control configuration.....	74
Database admin role-based access control configuration.....	75
NetApp ONTAP admin role-based access control configuration.....	76
IBM® admin role-based access control configuration.....	77
Pure Storage FlashArray admin role-based access control configuration.....	78
Dell PowerMax Storage admin role-based access control configuration.....	79
Configure tenants.....	81
Adding a tenant.....	81
Editing a tenant.....	82
Deleting a tenant.....	82
Best practices for configuring tenants.....	83
To assign resources to a tenant without granting the tenant users the ability to modify or delete the resources:.....	83
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 Defender Copy Data Management job definitions:.....	83
To assign resources to a tenant and allow the tenant users to create and run job sessions, reports and perform searches:.....	83
General recommendations.....	83
Identities.....	84
Identities overview.....	84
Adding a key.....	84
Adding an SSH key through the Generate a keypair for me method and registering an associated provider.....	84
Adding an SSH key through the I will provide a keypair method and registering an associated provider.....	85
Adding a credential.....	87
Editing a credential.....	87
Deleting a credential.....	87
Configure SLA policies.....	88
Creating an SLA policy.....	88
Configure IBM® Storage Virtualize SLA policies.....	89
Configure IBM® Storage Virtualize for Snapshot SLA policies.....	94
Configure NetApp ONTAP SLA policies.....	99
Configure Pure Storage FlashArray SLA policies.....	102
Configure Dell PowerMax Storage SLA policies.....	105
Configure Dell PowerFlex Storage SLA policies.....	107
Configure scripts.....	109
Uploading a script.....	109
Replacing a script.....	110
Deleting a script.....	110
Schedules.....	111
Creating a schedule.....	111
Editing a schedule.....	112
Deleting a schedule.....	112

Jobs	114
Jobs overview	114
Start, pause, and hold a job session	114
Starting a job session.....	114
Pausing and resuming a job session	115
Holding and releasing a job session	115
Canceling a job session	116
Monitor a job session	116
Monitoring a running job session.....	116
Filtering the list of jobs based on type or status.....	117
Viewing information about specific job sessions	117
Job definition overview	118
Job types	118
Inventory jobs.....	119
Creating a Database Inventory job definition	119
Creating a File System Inventory job definition.....	121
Creating an IBM® Storage Virtualize Inventory job definition	122
Creating an IBM® Storage Virtualize for Snapshot Inventory job definition	124
Creating a NetApp ONTAP Storage Inventory job definition	125
Creating a NetApp ONTAP File Inventory job definition	127
Creating a Pure Storage FlashArray Inventory job definition	130
Creating a VMware Inventory job definition	131
Creating a Dell PowerMax Storage Inventory job definition	133
Creating a Dell PowerFlex Storage Inventory job definition	134
Backup jobs	135
Creating an InterSystems Database Backup job definition	136
Creating an InterSystem IRIS application (on AIX®) scanning backup job.....	137
Creating an SAP HANA Backup job definition.....	139
Creating an Oracle Backup job definition.....	142
Creating an Oracle application (on AIX®) scanning backup job	146
Creating a SQL Backup job definition	150
Creating a File System Backup job	153
Creating an IBM® Storage Virtualize Backup job definition	156
Creating an IBM® Storage Virtualize for Snapshot Backup job definition.....	160
Creating a NetApp ONTAP Backup job definition.....	169
Creating a Pure Storage FlashArray Backup job definition.....	171
Creating a VMware Backup job definition	174
Creating a Dell PowerMax Storage backup job definition.....	178
Creating a Dell PowerFlex Storage backup job definition.....	179
Restore jobs.....	180
Creating an InterSystems Database Restore job definition	181
Creating an SAP HANA Restore job definition	185
Creating an Oracle Restore job definition	189
Creating a Microsoft™ SQL Restore job definition	195
Creating a File System Restore job definition.....	201
Creating an IBM® Storage Virtualize Restore job definition.....	204
Creating an IBM® Storage Virtualize for Snapshot Restore job definition	211
Creating a NetAPP ONTAP Restore job definition	216
Creating a Pure Storage FlashArray Restore job definition	223
Creating a VMware Restore job definition.....	227
Creating a Dell PowerMax Storage restore job definition	241
Creating a Dell PowerFlex Storage restore job definition	246
Renaming mount points and initialization parameter options.....	250
Using state and status arguments in postscripts	252
System jobs	253
Maintenance job.....	253
Creating a report job definition	253
Editing a job definition	254
Deleting a job definition	255
Searching	256
Searching overview	256
Searching for objects	256
Before you begin:	257
To search for objects:	257
To perform a basic search:.....	257
To perform an advanced search:	258
Wildcard considerations:.....	259
Viewing object details	260
Viewing NetApp ONTAP file details	261
Finding and restoring a file	262
Downloading search results	262
Browsing inventory.....	263
Browsing cataloged providers	263
Browsing through the Inventory on a previous date.....	264
Report.....	265

Report Overview	265
Running a report	265
Creating a customized report	266
Editing a customized report	267
Downloading a Report	267
Deleting a generated report	268
Application reports	269
Application configuration report	269
Application RPO Compliance report	271
System management reports	274
Catalog summary report	274
Configuration report	275
Job report	277
Job sessions report	278
File analytics reports	279
Quick view	280
Detail view	280
File usage by owner report	280
Files by age report	281
Files by category report	283
Files by size report	284
Protection compliance reports	285
File System RPO compliance report	286
IBM® Storage Virtualize RPO compliance report	288
NetApp ONTAP protection usage report	290
NetApp ONTAP RPO compliance report	292
Pure Storage FlashArray RPO compliance report	294
Recovery points report	296
Unprotected Virtual Machines report	297
VMware RPO compliance report	298
Storage protection reports	301
NetApp ONTAP OSSV relationship status report	302
NetApp ONTAP overprotected volumes report	302
NetApp ONTAP Qtree protection status report	303
NetApp ONTAP underprotected volumes report	304
NetApp ONTAP volume protection status report	305
NetApp ONTAP transition dependency report	307
Storage utilization reports	308
Quick View	309
Summary View	309
Detail View	309
IBM® Storage Virtualize Consistency Groups Report	310
IBM® Storage Virtualize Pools Report	311
IBM® Storage Virtualize Volumes Report	313
Instant Disk Restore Volumes Report	314
NetApp ONTAP Aggregates Report	315
NetApp ONTAP LUNs Report	316
NetApp ONTAP Orphaned LUNs Report	318
NetApp ONTAP Quotas Report	319
NetApp ONTAP Snapshots Report	320
NetApp ONTAP Volumes Report	322
Pure Storage FlashArray Volumes Report	323
Storage Capacity Report	325
VM and storage mapping report	326
VMware Datastores Report	326
VMware LUNs Report	328
VMware Orphaned Datastores Report	329
VMware Orphaned LUNs Report	330
VMware VM Snapshot Sprawl Report	332
VMware VM Sprawl Report	332
VMware VM Storage Report	334
Maintenance	336
Maintenance overview	336
Log on to the virtual appliance	336
Setting the time zone	337
Setting a Daylight Saving Time (DST) configuration	337
Collecting logs for troubleshooting	339
Collecting audit logs from the support menu	339
Collecting logs from the support menu	340
Collecting the IBM® Storage Defender Copy Data Management logs from the virtual appliance	340
Troubleshooting policy-based high availability and policy-based replication errors	341
Modifying job log options	342
Updating global settings	343
Managing the administrative console	344

Pre-upgrade resource adjustments to avoid out-of-memory issues and optimize performance	344
Upgrade IBM® Storage Defender Copy Data Management.....	346
Upgrading IBM® Storage Defender Copy Data Management appliance from 2.2.18 or later versions to 2.2.25 or later versions	346
Upgrading IBM® Storage Defender Copy Data Management appliance from 2.2.16 or 2.2.17 version to 2.2.18 or later versions	347
Upgrading IBM® Storage Defender Copy Data Management appliance from 2.2.15 or earlier versions to 2.2.16 or later versions	348
Backup and Restore the Catalog.....	350
Modifying the network settings.....	351
Uploading an SSL Certificate	352
Restoring a snapshot from a FlexGroup volume to another FlexGroup volume	353
Documentation and support	355
Documentation roadmap.....	355
Help System	355
User's Guide	355
About the help system	355
Starting Help	355
Before You Begin, Next Steps, and Related Topics.....	355
Search Help Feature	355
Security Management Topics.....	356
Reference topics	357
Search and filter guidelines	357
In search and filter fields:	357
Perform a basic search using inline search parameters:.....	357
Select, sort, and reorder columns.....	358
To select the columns to display:	358
To choose the column to sort on:	359
To change the order that the columns display:.....	359
LDAP username syntax.....	359
Return code reference	359
Tuning external parameterized configurations.....	361
Sentinel NFS tuning configuration	361
Sentinel scanning configuration	362
Remote script execution configuration	362
Clean rate configuration.....	362
In-backup mode configuration	362
Replication timeout configuration	364
Oracle agent code execution timeout configuration.....	364
The guestapps agent configuration properties.....	365
Latest blog posts	365
Frequently asked questions	367
Deployment	367
Resources	367
Connectivity	368
Cataloging	369
Operation	369
Backup/Restore Jobs.....	370
Oracle database support FAQ.....	371
Deployment and Registration.....	371
Oracle Backup creation workflow	371
Oracle Archive log management.....	373
Oracle RMAN integration	374
Pre and Post Scripts.....	374
Data Masking	374
Oracle Restore Workflow.....	375
System Requirements.....	378
Oracle Requirements	378
Microsoft™ SQL Server Support FAQ	379
What is SQL Server CDM? How does it help solve my challenges?	379
Deployment and Registration.....	379
SQL Server Backup workflow	379
SQL Server log management.....	381
Pre and Post Scripts.....	381
Data masking	381
SQL Server Restore Workflow	382
Self Service	384
System Requirements.....	384

SQL Support for VMware Virtual Machines.....	385
Registration and Authentication	385
SAP HANA HSR Cluster Server Support FAQ	387
What is SAP HANA System Replication (HSR Cluster)?	387
What steps should be taken in IBM® Storage Defender Copy Data Management after a primary node failover in the SAP HANA HSR Cluster?	387
What are the prerequisites for Instant Disk Restore (IDR)?	387
What are the prerequisites for Instant Database Restore (IDBR)?	388
Notices	389
Trademarks.....	390
Terms and conditions for product documentation	390
Privacy policy considerations	391
Glossary	392
Index	393

Note:

Before you use this information and the product it supports, read the information in “Notices” on page 389.

This edition applies to version 2, release 2, modification 28 of IBM® Storage Defender 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.

About this publication

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

Audience and purpose

This publication is intended for IBM® Storage Defender 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

New features and enhancements are available in IBM® Storage Defender Copy Data Management 2.2.28.

IBM® Storage Defender Copy Data Management enhancements

Feature	Description	Benefits
Support for Dell PowerFlex Storage arrays	IBM® Storage Defender Copy Data Management supports snapshot backup, replication, and restore on Dell PowerFlex Storage arrays.	Expands the support of applications by using Dell PowerFlex Storage arrays in the environment.
Support for policy-based high availability (PBHA) 3-site configuration for IBM Storage FlashSystem® version 9.1.0 and later	IBM® Storage Defender Copy Data Management supports PBHA 3-site configuration for IBM Storage FlashSystem® version 9.1.0 and later by using IBM Storage Virtualize Snapshot SLA policies.	Allows users to leverage the PBHA 3-site configuration by using IBM Storage FlashSystem®.
Support for IBM Storage FlashSystem® version 9.1.1	IBM® Storage Defender Copy Data Management supports IBM Storage FlashSystem® version 9.1.1	Expands the application support to the new features available in IBM Storage FlashSystem® version 9.1.1.
Support for Linux ext4 file system	IBM® Storage Defender Copy Data Management supports Linux ext4 file system.	Users can register, backup, and restore the Linux ext4 file system.
Support for IBM® Storage Defender Sentinel 1.1.12	IBM® Storage Defender Sentinel 1.1.12 is now qualified with IBM® Storage Defender Copy Data Management 2.2.28.	Use IBM® Storage Defender Sentinel 1.1.12 with IBM® Storage Defender Copy Data Management 2.2.28.

IBM® Storage Defender Sentinel

Feature	Description	Benefit
Application UI rebranding	The Home and Diagnostic and Reporting pages have been rebranded in the application UI.	Rebranding changes make the application consistent with user documentation and enhance the overall user experience.
Support for IBM® Storage Defender Copy Data Management 2.2.28	IBM® Storage Defender Sentinel 1.1.12 is now qualified with IBM® Storage Defender Copy Data Management 2.2.28.	Use IBM® Storage Defender Sentinel 1.1.12 with IBM® Storage Defender Copy Data Management 2.2.28.

Important: Review the [“Limitations and known issues” on page 11](#) topic before you start using the IBM® Storage Defender Copy Data Management application.

Limitations and known issues

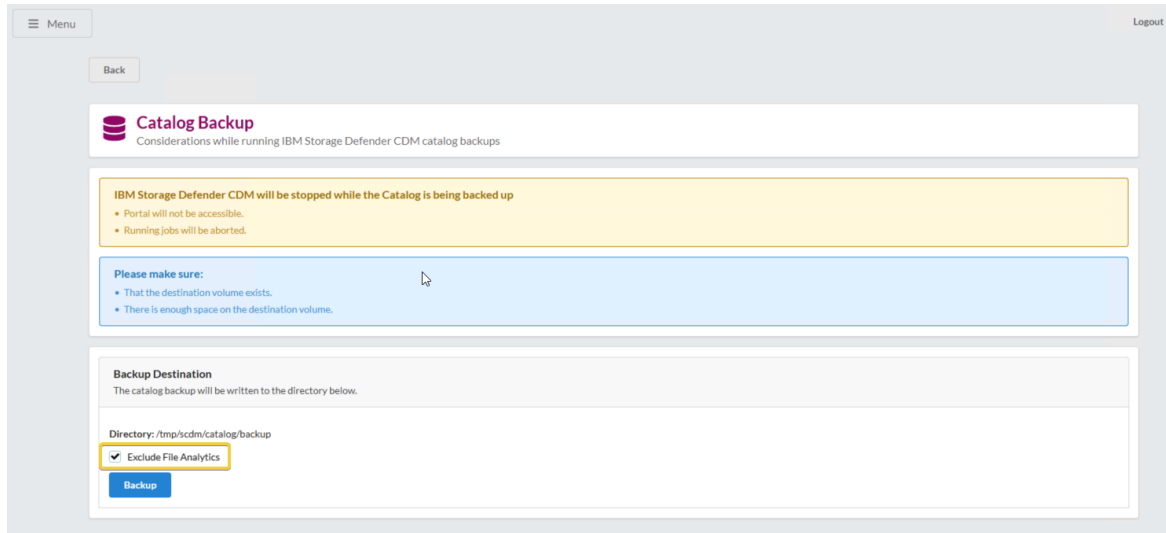
Understand the common limitations and known issues of the IBM® Storage Defender Copy Data Management application before you start using it.

Common limitations and known issues

The catalog backup procedure does not complete successfully when the **Exclude File Analytics** option is enabled

Description

When the user enables the **Exclude File Analytics** option and runs the catalog backup, the backup fails without generating an error message.



The screenshot shows the 'Catalog Backup' configuration page. At the top, there's a 'Menu' button and a 'Logout' link. Below that is a 'Back' button. The main section is titled 'Catalog Backup' with a subtitle 'Considerations while running IBM Storage Defender CDM catalog backups'. A yellow warning box states: 'IBM Storage Defender CDM will be stopped while the Catalog is being backed up'. Below this, two bullet points list considerations: 'Portal will not be accessible.' and 'Running jobs will be aborted.' A blue box with a cursor icon says 'Please make sure:' followed by two bullet points: 'That the destination volume exists.' and 'There is enough space on the destination volume.' The 'Backup Destination' section indicates the catalog backup will be written to the directory below. The 'Directory' is set to '/tmp/scdm/catalog/backup'. The 'Exclude File Analytics' checkbox is checked. A 'Backup' button is at the bottom.

Figure 1: Exclude File Analytics

Workaround

Run the catalog backup procedure without enabling the **Exclude File Analytics** option.

Catalog backup may not complete in an air-gapped environment

Description

In an air-gapped environment, catalog backup may not complete if an invalid DNS is configured by using the search option in the `/etc/resolv.conf` file.

Workaround

Complete the following steps:

1. On the Copy Data Management appliance, edit the `/etc/resolv.conf` file and remove the search option.
2. SSH to the Copy Data Management appliance with root access and issue the following commands:

```
systemctl restart k3s
kubectl delete pods -n kube-system --all
kubectl delete pods -n scdm --all
```

3. Login to **AdminConsole UI** and rerun the catalog backup.

The Virgo service may time out after an upgrade or reboot

Description

After upgrading or rebooting the Copy Data Management appliance, users may see the following screen when accessing the management GUI.




Figure 2: Accessing GUI error

Workaround

Restart the virgo service by using one of the following two methods:

- Login to **AdminConsole UI**, go to **Perform System Actions**, and restart the Copy Data Management appliance.
Or

SSH to the Copy Data Management appliance with root access and issue the following command:

```
systemctl restart virgo
```

After restarting the virgo service, wait for a few minutes before you access the management UI.

Running multiple backup or restore jobs in parallel for the same application is not supported

Description

Do not schedule or manually run multiple backup or restore jobs for the same application in parallel.

Workaround

None

After the application upgrade, the Admin UI connection may time out

Description

When you upgrade from a previous version to the latest version, you may see the following screen.




Figure 3: Upgrading Copy Data Management

Workaround

This may occur if the Admin UI connection times out. Ignore the message, refresh the Admin UI, and the appliance functions as expected..

Incremental FlashCopy is not supported for Local and Global Mirror backups of applications deployed virtually on VMs

Description

Using the incremental FlashCopy option within the IBM Storage Virtualize snapshot SLA can lead to issues when the application is deployed in a virtual environment. If the application is deployed virtually and has its data and logs on virtual disks (VMDKs), the incremental option does not work properly.

Workaround

Do not use FlashCopy Incremental for FlashCopy or Global Mirror in the SLA for backups of any applications (InterSystems, SAP HANA, Oracle, or SQL) on VMs. On VMs, the disks are considered as VMDKs.

Cancelling the job before inventory completion can lead to inconsistent or hung job state

Description

When executing a job pressing the cancel button before the inventory completion can lead to an inconsistent and hung state of the job.

Workaround

None

Upgrading to the latest version may take longer than previous upgrades

Description

Upgrade from the previous versions to the latest version can take longer than the previous upgrades.

Workaround

None

Do not use the same SLA policy for multiple backup jobs

Description

The use of the same SLA policy with multiple backup jobs could result in empty consistency groups on IBM Storage FlashSystem when the flash copies are condensed by the Copy Data backup job.

Workaround

Manually remove the empty flash copy mappings, along with the associated jobs and SLA policy. Then, create a new SLA policy and a backup job with one-to-one mapping between the SLA policy and the backup job.

Ignore warning message - Failed to add <DeviceName> to LVM devices file (!'Devices file not enabled.')

Description

When performing backup or restore operations for LVM-based applications, you may encounter the following message in the job log:

Failed to add <DeviceName> to LVM devices file (!'Devices file not enabled.')

This is just a warning message and does not affect the successful completion of the job.

Workaround

None

SAP HANA database specific limitations

For SAP HANA alternate hosts, the restore operation is not supported

Description

For SAP HANA alternate hosts, the restore (Instant Disk and Instant Database) operation is not supported. You need to restore it on the same source host.

Workaround

None

The Copy Data Management UI displays the "Identity cannot be deleted as it is in use" error when you try to delete both keys and credentials for an unregistered SAP HANA database

Description

If an SAP HANA database is registered with SSH keys and database credentials, unregistering it successfully still does not allow to remove the associated keys and database credentials. When you try to delete, the Copy Data Management UI displays the "Identity cannot be deleted as it is in use" error.

Workaround

Create new identity and private key for SAP HANA registration.

Oracle database specific limitations

The point-in-time recovery operation is not supported when datafiles are added between the recovery point and the previous backup

Description

The point-in-time recovery operation 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.

Workaround

None

The revert operation can only be used for backups that are created by using the IBM Storage Virtualize for Snapshot and Dell PowerMax Storage SLA policies

Description

For Oracle database applications, the revert operation is only available for the backups that are created by using the IBM Storage Virtualize for Snapshot and Dell PowerMax Storage SLA policies.

Workaround

None

Microsoft SQL Server specific limitations

Microsoft SQL database restore operations to a remote SQL server are not supported

Description

Microsoft SQL database restore operations to a remote SQL server are not supported due to an existing limitation of Microsoft SQL Server.

Workaround

None

The revert operation is not supported for SQL Server Failover Clustering and AlwaysOn configuration for Microsoft SQL databases registered as virtual in the Copy Data Management appliance

Description

The revert operation is not supported for SQL Server Failover Clustering and AlwaysOn configuration for Microsoft SQL databases registered as virtual in the Copy Data Management appliance.

Workaround

None

VMware specific limitations

The restore operation to an older snapshot is not supported if a new snapshot is used for mirror replication

Description

You cannot restore to an older snapshot if you have used a new snapshot for mirror replication.

Workaround

None

The revert operation is not supported for individual VMs of VMware hypervisor

Description

The revert operation is available only for datastores and not supported for individual VMs of VMware hypervisor.

Workaround

None

Dell PowerMax Storage specific limitations

The revert operation from a remote copy is not supported

Description

For Dell PowerMax Storage arrays, the revert operation from a remote copy is not supported.

Workaround

If the user selects a remote copy from the available options for the revert operation, the restore job will fail. No workaround is available for this limitation.

Maintenance jobs may fail to delete snapshots from Dell PowerMax Storage arrays

Description

For Dell PowerMax Storage, maintenance jobs may fail to delete snapshots from the storage arrays as expected after the retention period ends if the source volumes (i.e., volumes for which snapshots are taken) have identical names across multiple arrays managed under the same Unisphere instance. The naming conflict may result in incorrect volume resolution during snapshot deletion, causing the cleanup process to fail.

Workaround

Manually delete the snapshots after their retention period ends.

Note: If you reattempt a revert by using a local copy after previously attempting a revert from a remote copy (which is not supported), subsequent attempts to perform a revert may fail.

Dell PowerFlex Storage specific limitations

The revert operation from a remote copy is not supported

Description

For Dell PowerFlex Storage arrays, the revert operation from a remote copy is not supported.

Workaround

If the user selects a remote copy from the available options for the revert operation, the restore job will fail. No workaround is available for this limitation.

Remote snapshot and restore require higher privileges in Dell PowerFlex due to a known limitation in PowerFlex Manager 4.6.1

Description

You can successfully execute remote snapshot and restore workflows only when the associated Dell PowerFlex Storage arrays are registered with super-user-level Dell PowerFlex credentials in the Copy Data Management appliance.

Workaround

None

IBM Storage Virtualize for Snapshot specific limitations

Recovery are points not visible for restore when mixed retention policies exist on same volume

Description

When two backup jobs are created for the same volume, one by using retention by count and the other by using retention by days, and neither job is scheduled, with only one backup each, an issue can occur. If the day-based retention job creates its backup after the count-based job, then once the day-based recovery points expire, the recovery points for the count-based job does not appear when creating a restore job, even though they still exist in the catalog. The issue does not occur if either job is scheduled and has at least one unexpired recovery point.

Workaround

Execute the backup job with retention by day and take a new backup to ensure that recovery points for the other backup job are visible when creating a restore job.

Incomplete Application RPO compliance report

Description

You may see an incomplete or invalid Application RPO compliance report for IBM Storage Virtualize for Snapshot.

Workaround

None

Policy-based high availability (PBHA) specific limitations

Creating a volume clone from snapshot is not allowed in a PBHA partition

Description

As per the FlashSystem design, creating a volume clone from a snapshot is not permitted within the PBHA partition. Therefore, you cannot restore or create copies of volumes or databases on the original host or server by using PBHA backups. Any backups taken by using PBHA policy cannot be used to perform disk or database restore on the original host. However, these backups can be used to revert the original volume on Active Management Site (AMS) or to create a clone in a different partition or outside the PBHA partition.

Workaround

This behavior is by design and not considered a limitation.

Any create or delete operation is supported only on the AMS node

Description

For PBHA, any create or delete operation must be performed on the AMS node. Revert of volume from a snapshot is only allowed on the current AMS node. Therefore, to revert a volume from a snapshot, you must choose the backup that is created on the node which is currently acting as AMS.

PBHA backups are first cataloged on the primary site, then on secondary site

Description

For PBHA backup, first snapshots are taken and cataloged on the primary site. Subsequently, snapshots are cataloged on the secondary site. Any errors encountered with failed or partially failed volumes are recorded.

Workaround

None

For PBHA backup of storage volumes, the consistency group option is ignored

Description

For PBHA storage volume backups, the consistency group option is disregarded. If the volumes belong to different volume groups, they are grouped by their respective volume groups. Individual snapshots are taken for each volume group.

Workaround

None

All selected databases must be in the same volume group for application backup

Description

For backup of application (database) volumes, all selected databases must be in the same volume group.

Workaround

None

The database restore operation on a SQL database from a remote copy is not supported

Description

Restoring a SQL database from a remote copy is not supported when the backup is created by using a PBHA policy. Only local disk-based restores are permitted. Even if the database resides on a host with access to both primary and disaster recovery (DR) site storage, the snapshot is taken at the DR site is considered a remote copy for the host.

Workaround

This limitation is due to an existing Microsoft SQL database limitation. No workaround is available for this limitation. You can use the revert volume functionality in case you want to restore the database. The revert volume functionality replaces the original volume with the data from the selected snapshot.

Database restore for SAP is not supported on original as well as alternate server when backup is created by using a PBHA policy

Description

When a backup is created by using a PBHA policy, you cannot use it to do database restore. Restore cannot be performed on either the original or alternate server.

Workaround

This limitation is due to IBM FlashSystem design. You can use the revert volume functionality in case you want to restore the database. The revert volume functionality replaces the original volume with the data from the selected snapshot.

Policy-based replication (PBR) specific limitations

Remote array snapshots for PBR are crash-consistent only

Description

Remote array snapshots for PBR are crash-consistent only, whereas snapshots taken on the primary site are app-consistent.

Workaround

None

Consistency group requires unified grouping

Description

If consistency group is enabled for PBR backup of storage volumes, all volumes must be in the same volume group.

Workaround

None

Creating a single SLA policy with two sub-policies is not supported

Description

If you need snapshots at both the primary and disaster recovery (DR) sites in PBR setup, you cannot create an SLA policy with two sub-policies, one for primary site and another for DR site.

Workaround

When setting up a backup job in a PBR setup, if you need snapshots at both the primary and DR sites, you must create two separate SLA policies, one for the primary site, and another for the DR site.

The database restore operation on a SQL database from a remote copy is not supported

Description

Restoring a SQL database from a remote copy is not supported when the backup is created using a PBR policy, which is the only available option in a PBR setup. Only local disk-based restores are permitted. Even if the database resides on a host with access to both primary and DR) site, the snapshot is taken at the DR site is considered a remote copy from the host.

Workaround

This limitation is due to an existing Microsoft SQL database limitation. No workaround is available for this limitation.

Sentinel specific limitations

Limitations in Scanning VMs with VMware Snapshots on ESXi 8.0.3 Using Sentinel

Description

Sentinel server with version 8.10 or higher can scan virtual machines on ESXi version 8.0.3. However, if those VMs have VMware snapshots, the host will not scan them. Sentinel server with version 8.9 or earlier can also scan virtual machines, but they cannot scan any VMs that have VMware snapshots, regardless of the ESXi version.

Workaround

None

Orphaned volumes on AIX proxy server with Sentinel

Description

During the scanning of AIX-based workloads, the file system unmount operation on the AIX proxy server may exceed the expected time limit and fail to complete.

Workaround

As a workaround, you can follow one of the following methods, either on the Copy Data Management appliance or the AIX proxy server:

- Copy Data Management side workaround:

- a. Manually clean up the stale volume groups from the AIX proxy server.
 - b. Increase the unmount command timeout in the `protection.application.AIX.proxy.unmount.timeout.seconds` system property, which is available in the `com.syncsort.dp.xsb.serviceprovider.properties` file.
- AIX side workaround:
 - Coordinate with the AIX team to identify the AIX version where the unmount command hang issue has been fixed.

IBM® Storage Defender Copy Data Management overview

IBM® Storage Defender 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 and DevOps. 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 Defender Copy Data Management delivers “in-place” copy data management to enterprise storage arrays from IBM®, Dell PowerMax Storage, 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 Defender 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 Defender Copy Data Management, see [Requirements](#).

Getting off to a quick start

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

Related information

[Start IBM Storage Defender Copy Data Management](#)

[Sites and providers overview](#)

[Start, pause, and hold a job session](#)

[Searching for objects](#)

[Report Overview](#)

Start IBM® Storage Defender Copy Data Management

Start IBM® Storage Defender Copy Data Management to use the application and its features.

Procedure

To Start IBM® Storage Defender 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 Defender Copy Data Management.

2. Enter your username and password, which is provided by the IBM® Storage Defender Copy Data Management System Administrator.
If this is your first time logging on to IBM® Storage Defender 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®, Dell PowerMax Storage, Pure 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.

Related information

[Registering a Dell PowerMax provider](#)

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.

Related information

[Registering a storage provider](#)

Run a job

A job that is based on an Inventory job definition discovers object information, catalogs it, and populates the IBM® Storage Defender 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 Defender 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.

User administration and security management

IBM® Storage Defender Copy Data Management provides users the opportunity to rapidly locate files and objects on IBM®, Dell PowerMax Storage, NetApp, and Pure Storage FlashArray devices along with VMware ESXi, Oracle, SQL, SAP HANA, and InterSystems hosts. IBM® Storage Defender 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 Defender 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 Defender 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 Defender 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 Defender Copy Data Management functionality.

Sensitive data is encrypted when stored.

Data in transit is also protected. IBM® Storage Defender 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 Defender 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 Defender Copy Data Management identifies the following types of sensitive data: native user credentials, IBM®, Dell PowerMax Storage, 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 Defender 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 Defender Copy Data Management to use LDAP, and configuring IBM® Storage Defender Copy Data Management to use HTTPS.

Following are the security management functions in IBM® Storage Defender 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 Defender 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 Defender 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 Defender Copy Data Management database. Obtaining and storing this sensitive data constitutes the basic function of the IBM® Storage Defender Copy Data Management application. This data is subject to the user data security requirements.

Ports

The following ports are used by IBM® Storage Defender Copy Data Management:

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 Defender Copy Data Management
68	bootpc in DHCP clients, DHCP Listener UDP	Service used by IBM® Storage Defender Copy Data Management
80/443	http/https	Service used by IBM® Storage Defender Copy Data Management
389	LDAP, non-TLS connection for Lightweight Directory Access Protocol	Service used by IBM® Storage Defender Copy Data Management
443	smtp, TLS connection for Simple Mail Transfer Protocol	Service used by IBM® Storage Defender Copy Data Management
636	LDAP, TLS connection for Lightweight Directory Access Protocol	Service used by IBM® Storage Defender Copy Data Management
1433	sql, SQL Service	Service used by IBM® Storage Defender Copy Data Management
4369	epmd, Erlang port mapper	Service used by IBM® Storage Defender Copy Data Management
5480	ssl/http, vami	Port open within the firewall, IBM® Storage Defender Copy Data Management 2.2.5 and earlier
5985	smtp, TLS connection for Simple Mail Transfer Protocol	WinRM, Windows™ Remote Management

Port	Service	Comment
8090	admin console, IBM® Storage Defender Copy Data Management Administrative Console	Port open within the firewall
8092	adminconsole, IBM® Storage Defender Copy Data Management Administrative Console	Port open within the firewall, IBM® Storage Defender 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 Defender Copy Data Management services
9090	Liberty Server	Service used to serve the Knowledge Center documentation
27017	MongoDB mongod	Service used by IBM® Storage Defender Copy Data Management
55672	rabbitMQ, RabbitMQ administrative	Service used by IBM® Storage Defender Copy Data Management

Related information

[Registering a storage provider](#)

[Role-based access control overview](#)

Installation and setup

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

Deployment checklist

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

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

Table 4: Pre-Deployment checklist

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

Table 5: Deployment checklist

Step	Action	Related Topic	
1	Install IBM® Storage Defender Copy Data Management by deploying an OVF template to create a virtual appliance containing the application on a VMware host.	Installing IBM® Storage Defender Copy Data Management as a virtual appliance	
2	To access IBM® Storage Defender Copy Data Management, upload a valid product key to the virtual machine where IBM® Storage Defender Copy Data Management is deployed.	Installing IBM® Storage Defender Copy Data Management as a virtual appliance	
3	Launch IBM® Storage Defender Copy Data Management to set a new Super User password.	Start IBM® Storage Defender 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	

Table 6: Post-Deployment checklist

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

Interface	Location	Default Credentials	Notes®
IBM® Storage Defender Copy Data Management User Interface	<code>https://<HOSTNAME_OR_IP>:8443/portal/</code>	Username: admin Password: password	This is the main IBM® Storage Defender 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 Defender Copy Data Management users.

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

Interface	Location	Default Credentials	Notes®
IBM® Storage Defender 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 Defender 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 Defender Copy Data Management OVA.

VMware vSphere™ Privileges

In IBM® Storage Defender 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
Distributed switch	<ul style="list-style-type: none"> • Port configuration operation • Port setting operation

vCenter Server Object	Required Privileges
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)

vCenter Server Object	Required Privileges
Virtual Machine > Guest Operations	<ul style="list-style-type: none"> • Guest Operation Modifications • Guest Operation Program Execution • Guest Operation Queries
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 Defender Copy Data Management as a virtual appliance

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

Procedure

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

1. Use the vSphere Client to deploy IBM® Storage Defender Copy Data Management. Go to **Development > Actions** menu, choose **Deploy OVF Template**. The Deploy OVF template window opens.

2. Enter the URL for the IBM® Storage Defender 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 Defender 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 Defender Copy Data Management is valid. Click **Ignore** to close the warning in the vCenter OVA UI wizard when you deploy IBM® Storage Defender 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 Defender 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 337](#).

- After the first use, you can enable additional users to logon by linking to an LDAP server.
- Start IBM® Storage Defender Copy Data Management and begin using it from any supported web browser. See [“Start IBM Storage Defender Copy Data Management” on page 21.](#)

Related information

[Start IBM Storage Defender Copy Data Management](#)

Installing IBM® Storage Defender Sentinel Security Scan

The topic section cover installing IBM® Storage Defender Sentinel Security Scan.

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

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

Attention: A federation consists of two or more IBM® Storage Defender Sentinel Security Scan software systems together to share work. Install the IBM® Storage Defender Sentinel Security Scan on the federation manager machines and then add a license to the manager. Now you can install new member IBM® Storage Defender 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 Defender 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 Defender 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 Defender 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 Defender Copy Data Management"
```

Result

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

Starting IBM® Storage Defender Copy Data Management

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

Before you begin

- IBM® Storage Defender Copy Data Management must be installed before starting the application. See [“Installing IBM Storage Defender Copy Data Management as a virtual appliance” on page 32](#).
- The System Administrator must provide you with the IP address for the virtual appliance and the IBM® Storage Defender Copy Data Management username and password.

Procedure

To start IBM® Storage Defender 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 Defender Copy Data Management.

2. In the logon dialog, enter your username and password, which is provided by the IBM® Storage Defender Copy Data Management System Administrator.
If this is your first time logging on to IBM® Storage Defender 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 Defender 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 68](#).
- Add storage systems and virtual machine resources to the IBM® Storage Defender Copy Data Management database. See [“Registering a storage provider” on page 40](#) and [“Jobs overview” on page 114](#).
- Search or browse for objects that match certain criteria. See [“Searching for objects” on page 256](#) and [“Browsing inventory” on page 263](#).
- Generate reports with predefined or customized parameters. See [“Report Overview” on page 265](#).

Related information

[Installing IBM Storage Defender Copy Data Management as a virtual appliance](#)

Dashboard

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

Before you begin, review IBM® Storage Defender Copy Data Management features and data flow information. See [IBM® Storage Defender 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 Defender 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 Defender 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 information

[Sites and providers overview](#)

[Role-based access control overview](#)

[Jobs overview](#)
[Monitor a job session](#)
[Searching overview](#)
[Report Overview](#)

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 Defender 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 Defender Copy Data Management. Unlike storage systems and VMware servers, SMTP servers are not cataloged.
- **VMware servers:** Supported types includes vCenter and ESX/ESXi hosts.
- **Dell PowerMax Storage system:** Supported storage platforms include Dell PowerMax FlashArray.

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 Defender 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 information

[Registering a storage provider](#)

[Viewing a provider](#)

[Editing a provider](#)

[Unregistering a provider](#)

Adding a site

Add sites to define a grouping of providers based on their location in your IBM® Storage Defender Copy Data Management environment. Once sites are created in IBM® Storage Defender 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 38](#) 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 storage provider” on page 40](#) and [“Editing a provider” on page 61](#).

Related information

[Editing a site](#)

[Deleting a site](#)

Editing a site

Revise site names and descriptions to reflect location changes in your IBM® Storage Defender 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 storage provider” on page 40](#) and [“Editing a provider” on page 61](#).

Related information

[Adding a site](#)
[Deleting a site](#)

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 storage provider” on page 40](#) and [“Editing a provider” on page 61](#).

Related information

[Adding a site](#)
[Editing a site](#)

Registering a storage provider

Providers are physical servers that host objects and attributes. Once a provider is registered in IBM® Storage Defender 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 Defender 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 Defender 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 39](#).

- 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®.

Note: Supported filesystem types:

- AIX - JFS2
- Windows - NTFS
- Linux - xfs, ext3, and ext4

System Credential

Select or create your file system or volume's credentials. See [“Identities overview” on page 84](#).

6. Click **OK**. IBM® Storage Defender 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 Defender 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 39](#).
- 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 Defender 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 Defender 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 84](#).

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 84](#).

6. Click **OK**. IBM® Storage Defender 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 Defender Copy Data Management appliance and the server, and installs an IBM® Storage Defender 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 39](#).
- 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 Defender 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 Defender 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 84](#).

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 Defender 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 Defender Copy Data Management appliance and the server, and installs an IBM® Storage Defender 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 39](#).
- 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 an 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.

When registering an SAP HANA System Replication (SAP HSR) cluster, use the cluster’s virtual IP address or host name. Additionally, ensure that all participating nodes in the cluster are registered.

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 Defender 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.

Authentication

Select or create your SAP HANA operating system and database credentials. See [“Identities overview” on page 84](#) if you want to define users before registering the provider. For authentication, you can select one of the following authentication options:

Key

Provide the username in the **Username** field. You can select an existing key or create a new key. To select an existing key, click **Select**, or click **New** to create a new key.

Password

Select an existing credential or create a credential in the **System Credential** section. To select an existing credential, click **Select**, or click **New** to create a new credential.

System Credentials

Select or create your SAP HANA and database host operating system user’s credentials. The operating system user must have password less sudo privileges to execute commands as sXXadm user. For more information about sudo options, see [SAP HANA requirements](#).

Key

IBM® Storage Defender Copy Data Management supports the SSH key-based operating system authentication for SAP HANA database servers. Allows users to authenticate the SAP HANA application host by using the SSH key.

Important: The SAP HANA registration with IBM® Storage Defender Copy Data Management only works when the SSH key pair is generated by using the `ssh-keygen -t rsa -m PEM` command.

Database Credentials

The database credentials can be either for SYSTEM user or normal user that exists in SYSTEMDB and SXX tenant database with same username, password, and appropriate permissions. For more information on creating user by using HDBSQL command line interface, see [SAP HANA requirements](#).

6. Click **OK**. IBM® Storage Defender 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 Defender Copy Data Management appliance and the server, and installs an IBM® Storage Defender 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 39](#).
- 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 Defender Copy Data Management appliance must be resolvable and route-able by the DNS server; the physical SQL server will communicate back to IBM® Storage Defender Copy Data Management through DNS.
 - Note that for physical SQL servers, you must allow outgoing connections to port 8443 on the IBM® Storage Defender 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 Defender 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 Defender Copy Data Management.

Run Inventory job after registration

If selected, IBM® Storage Defender 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 84](#).

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 Defender 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 Defender Copy Data Management appliance and the server, and installs an IBM® Storage Defender 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 39](#).
- 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 Defender 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 84](#).

Note: If upgrading from a previous version of IBM® Storage Defender 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 Defender 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 Defender Copy Data Management.

Registering an IBM Storage Virtualize for Snapshot 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 39](#).
- 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 for Snapshot 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 for Snapshot**.
3. Right-click **IBM Storage Virtualize for Snapshot**, then click **Register**. The **Register** dialog opens.
4. Populate the fields in the dialog:

Site

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

Comment

Optional provider description.

Run Inventory job after registration

If selected, IBM® Storage Defender 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 for Snapshot credentials. See [“Identities overview” on page 84](#).

Note: If upgrading from a previous version of IBM® Storage Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 39](#).
- 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 Defender 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 84](#).

Note: If upgrading from a previous version of IBM® Storage Defender 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 Defender 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 39](#).
- 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 Defender 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 84](#).

Note: If upgrading from a previous version of IBM® Storage Defender 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 Defender 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 39](#).
- 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 84](#).

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

This procedure describes how to register a Security Scan Server in IBM® Storage Defender 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 44](#).

Restriction:

- When a Security Scan Server is added to IBM® Storage Defender 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 Defender 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 Defender Sentinel Security Scan software with the license. For more information, refer to [“Installing IBM Storage Defender Sentinel Security Scan” on page 34](#).

Procedure

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

1. Log in to IBM® Storage Defender 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.

Result

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 39](#).
- 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 Defender 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 84](#).

Note: If upgrading from a previous version of IBM® Storage Defender 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.

Comment

Optional provider description.

5. Click **OK**. IBM® Storage Defender 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 39](#).
- 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 84](#).

Note: If upgrading from a previous version of IBM® Storage Defender 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.

Comment

Optional provider description.

Run Inventory job after registration

If selected, IBM® Storage Defender 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 Defender 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 63](#).

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 39](#).
- 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 Defender Copy Data Management creates a high-level Inventory job and automatically catalogs the objects on the provider.
7. Click **Register**. IBM® Storage Defender 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 63](#).

What to do next

- Once providers are available in IBM® Storage Defender Copy Data Management and associated with a site, assign them to a resource pool. See [“Configure resource pools” on page 69](#).
- Add credentials to virtual machines in a VMware environment. See [“Adding credentials to a virtual machine” on page 63](#).

Related information

[Sites and providers overview](#)

[Viewing a provider](#)

[Editing a provider](#)

[Unregistering a provider](#)

[Uploading an SSL Certificate](#)

[LDAP username syntax](#)

Registering a Dell PowerMax provider

Before you begin

- Create a site to assign to your provider. A site is a user-defined grouping of providers that is based on location. See [“Adding a site” on page 39](#).
- When you register providers, if possible, assign all related sources, such as hosting vCenters and related storage systems, to the same site to manage your storage environment efficiently.
- For application server, if possible, keep all resources, such as hosting vCenters and related storage systems, which are configured in the same site to manage your storage environment efficiently.

Procedure

To register a Dell PowerMax provider, complete the following steps:

1. Click the **Configure** tab. On the **Views** page, select **Sites & Providers**, then select the **Providers** tab. The **Provider Browser** page is displayed.
2. On the **Provider Browser** page, select **Dell PowerMax**.
3. Right-click **Dell PowerMax**, then click **Register**. The **Register Dell PowerMax** dialog opens with the **Enter details** page.
4. Populate the following fields:
Site

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

Unisphere alias

A user-defined name for Dell Unisphere system. The value can be the same as the hostname or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique.

Unisphere IP address / FQDN

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

Port

To register Dell PowerMax, the default communication port is 8443. The default port number is auto populated in the **Register PowerMax Storage** window. Enter a different port number if you do not want to use the default port number.

Comment

Add a comment.

Run Inventory job after registration

If you select, IBM® Storage Defender Copy Data Management creates a high-level inventory job and automatically catalogs the objects on the provider. The Inventory job can take longer time to complete.

Credentials

Select or create your Dell Unisphere credentials. See [“Identities overview” on page 84](#).

5. Click **Next**. The **Select arrays** page appears. On the **Select arrays** page, select Dell PowerMax Storage providers that are registered to the Dell Unisphere system.
 - a. On the **Select arrays** page, select the arrays that you want to register by selecting the checkboxes in the table.

If a message, the connection is unsuccessful, is displayed, review your entries. If your entries are correct and the connection is unsuccessful, contact your system administrator.

Note: If the selected Dell Unisphere does not have any Dell PowerMax Storage systems available, register the provider by using a different Dell Unisphere instance.

Note: On the **Select arrays** page, all the arrays that are available in Dell Unisphere are displayed. All the arrays are selected by default. You can select or deselect specific arrays.

6. Click **Register**. The **Registration summary** page is displayed. The summary contains the following details:
 - The status of the Dell PowerMax registration
 - Registration errors (if any).
7. Click **Close** to complete the registration.

Registering a Dell PowerFlex Storage provider

Before you begin

- Create a site to assign to your provider. A site is a user-defined grouping of providers that is based on location. See [“Adding a site” on page 39](#).
- When you register providers, if possible, assign all related sources, such as hosting vCenters and related storage systems, to the same site to manage your storage environment efficiently.
- For the application server, if possible, keep all resources, such as hosting vCenters and related storage systems, which are configured in the same site to manage your storage environment efficiently.

Procedure

To register a Dell PowerFlex Storage provider, complete the following steps:

1. Click the **Configure** tab. On the **Views** page, select **Sites & Providers**, then select the **Providers** tab. The **Provider Browser** page is displayed.
2. On the **Provider Browser** page, select **Dell PowerFlex Storage**.
3. Right-click **Dell PowerFlex Storage**, then click **Register**. The **Register Dell PowerFlex Storage** dialog opens with the **Enter details** page.
4. Populate the following fields:

Site

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

PowerFlex Manager alias

A user-defined name for Dell PowerFlex Manager system. The value can be the same as the hostname or it can be a meaningful name that is used within your organization to refer to the provider. Provider names must be unique. The Copy Data Management appliance appends the Dell PowerFlex Storage system ID to the specified alias as a suffix. You can rename the alias later by editing the corresponding provider, if required.

PowerFlex Manager IP / FQDN

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

Port

To register Dell PowerFlex Storage, the default communication port is 443. The default port number is auto populated in the **Register PowerFlex Storage** window. Enter a different port number if you do not want to use the default port number.

Comment

Add a comment.

Run Inventory job after registration

If you select, IBM® Storage Defender Copy Data Management creates a high-level inventory job, execute it and automatically catalogs the objects on the provider. The Inventory job may take longer time to complete depending on the number of objects on the storage system.

If you do not select this option, IBM® Storage Defender Copy Data Management creates a default inventory job but does not execute it. You must manually run the job from the **Jobs** section.

Credentials

Select or create your Dell Unisphere credentials. See [“Identities overview” on page 84](#).

5. Click **Next**. The **Select arrays** page appears. On the **Select arrays** page, select Dell PowerFlex Storage arrays that are associated to the PowerFlex Manager system.
 - a. On the **Select arrays** page, select the arrays that you want to register by selecting the checkboxes in the table.

If a message, the connection is unsuccessful, is displayed, review your entries. If your entries are correct and the connection is unsuccessful, contact your system administrator.

Note: If the selected PowerFlex Manager does not have any Dell PowerFlex Storage systems available, register the provider by using a different PowerFlex Manager instance where the arrays are available.

Note: On the **Select arrays** page, all the arrays that are available in PowerFlex Manager are displayed. All the arrays are selected by default. You can select or deselect specific arrays.

6. Click **Register**. The **Registration summary** page is displayed. The summary contains the following details:
 - The status of the Dell PowerFlex Storage registration.
 - Registration errors (if any).

7. Click **Close** to complete the registration.

Viewing a provider

Navigate through the Provider Browser to view a list of providers 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. Select 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 data centers configured on the vCenter. Select the **Datacenters** tab to view the list of **Datastores**, **Hosts**, and **VMs** filtered by the selected data center.

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

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.

Dell PowerMax Storage providers

Select a Dell PowerMax provider.

The **General** tab provides a summary of site, host name, IP address, name of the array, model, microcode, and other useful information.

The **Volumes** tab provides a list of all the volumes that are part of all the Dell PowerMax Storage arrays being managed by the Dell Unisphere host.

The **Volume Groups** tab provides a list of all the storage groups that are present on the Dell PowerMax Storage arrays that are managed by the Dell Unisphere host.

The **Hosts** and **Host Groups** tabs provide a list of storage hosts that use the Dell PowerMax Storage LUN resources of the storage system.

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**.

Dell PowerFlex Storage providers

Select a Dell PowerFlex Storage provider.

The **General** tab provides a summary of site, host name, IP address, name of the array, model, and other useful information.

The **Volumes** tab provides a list of all the volumes that are part of all the Dell PowerFlex Storage systems.

The **Hosts** tabs provides a list of storage hosts that are available on the Dell PowerFlex Storage systems.

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 storage provider” on page 40](#).

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 on storage providers that are recently cataloged. See [“Jobs overview” on page 114](#).

Related information

[Sites and providers overview](#)

[Registering a storage provider](#)

[Editing a provider](#)

[Unregistering a provider](#)

Editing a provider

Edit 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 84](#).

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 114](#).

Related information

[Sites and providers overview](#)

[Registering a storage provider](#)

[Viewing a provider](#)

[Unregistering a provider](#)

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 61](#).

Restriction: A provider cannot be deleted if it is assigned to a Resource Pool. Remove your providers from Resource Pools before deleting.
Before unregistering a provider, delete all SLA policies and jobs associated with the provider and execute the maintenance job.

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 Defender 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 information

[Sites and providers overview](#)

[Registering a storage provider](#)

[Viewing a provider](#)

[Editing a provider](#)

Adding credentials to a virtual machine

Some features in IBM® Storage Defender 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 Defender Copy Data Management. See [“Registering a storage provider”](#) on page 40, and [“Creating a Database Inventory job definition”](#) on page 119.

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 information

[Registering a storage provider](#)

[Creating a Database Inventory job definition](#)

Adding credentials for multiple virtual machines

Before you begin

At least one VMware provider must be registered and cataloged in IBM® Storage Defender Copy Data Management. See [“Registering a storage provider” on page 40](#), and [“Creating a Database Inventory job definition” on page 119](#).

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 information

[Registering a storage provider](#)

[Creating a Database Inventory job definition](#)

Requiring multifactor authentication

Starting with IBM® Storage Defender Copy Data Management 2.2.21, you can set up multifactor authentication (MFA) on IBM® Storage Defender 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 the time-based one-time (TOTP) for a user account, you must manually enable it in the IBM® Storage Defender 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 Defender Copy Data Management server time. You must use NTP server to ensure the time sync between IBM® Storage Defender 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 Defender Copy Data Management ADMIN account.

Procedure

To enable time-based one-time (TOTP) multifactor authentication on the IBM® Storage Defender 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 Defender 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 Defender Copy Data Management.

Procedure

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

1. Sign on to the IBM® Storage Defender Copy Data Management using the user ID and password of MFA-enabled user. The IBM® Storage Defender 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 Defender 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 Defender 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 Defender 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.

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 Defender Copy Data Management accounts.

Through role-based access control you can tailor IBM® Storage Defender 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 Defender 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 Defender 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 Defender 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 Defender Copy Data Management, such as storage devices and LDAP servers, can be included in a resource pool, along with individual IBM® Storage Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender Copy Data Management and use its functions, you must first add the user to IBM® Storage Defender 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 information

[Configure resource pools](#)

[Configure roles](#)

[Configure accounts](#)

[VMware admin role-based access control configuration](#)

[NetApp ONTAP admin role-based access control configuration](#)

[IBM admin role-based access control configuration](#)

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 Defender Copy Data Management, such as storage devices and LDAP servers, can be included in a resource pool, along with individual IBM® Storage Defender 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 Defender 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 [Adding a site](#).
- Add providers to IBM® Storage Defender Copy Data Management and associate them with a site. See [Registering 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 Defender 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 Defender Copy Data Management features. Updated resource pool settings take affect once accounts associated with the resource pool log in to IBM® Storage Defender 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 Defender 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 resource pools](#).

Related information

[Configure roles](#)

[Configure accounts](#)

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

Issue

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 Defender 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 effect once accounts associated with the role log in to IBM® Storage Defender Copy Data Management.

About this task

Revise a role to change the resources and permissions assigned to the role. Updated role settings take effect once accounts associated with the role log in to IBM® Storage Defender 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 role 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 Defender Copy Data Management or can be imported as an LDAP group. See [Configure Accounts](#).

Related information

[Configure resource pools](#)

[Configure accounts](#)

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 Defender Copy Data Management and use its functions, you must first add the user to IBM® Storage Defender 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 [Adding a site](#).
- Add providers to IBM® Storage Defender Copy Data Management and associate them with a site. See [Registering a provider](#).
- Once providers are available in IBM® Storage Defender 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 Defender 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 Defender Copy Data Management

You can import LDAP groups into IBM® Storage Defender 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 affect once the account logs in to IBM® Storage Defender 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 Defender 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 Defender Copy Data Management resources as well as the necessary permissions to interact with the resources. See [Configure resource pools](#) and [Configure Roles](#).

Related information

[Configure resource pools](#)

[Configure roles](#)

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

Table 9: 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 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.
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 information

[Configure resource pools](#)

[Configure roles](#)

[Configure accounts](#)

Database admin role-based access control configuration

The Database admin (DBA) role is applicable for any database that IBM® Storage Defender 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 Defender Copy Data Management users to run Backup and Restore jobs by using database resources without having direct access to them.

Resource pool configuration

Table 10: Resource pool configuration	
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.
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 Defender Copy Data Management users to run Backup and Restore jobs using NetApp ONTAP resources without having direct access to them.

Resource pool configuration

Table 11: 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.

Resource pool configuration	Steps
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.
NetApp 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 information

[Configure resource pools](#)

[Configure roles](#)

[Configure accounts](#)

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 Defender Copy Data Management users to run Backup and Restore jobs using IBM® resources without having direct access to them.

Resource pool configuration

Table 12: 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.

Resource pool configuration	Steps
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 information

[Configure resource pools](#)

[Configure roles](#)

[Configure accounts](#)

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 Defender Copy Data Management users to run Backup and Restore jobs using Pure Storage resources without having direct access to them.

Resource pool configuration

Table 13: 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.

Resource pool configuration	Steps
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 information

[Configure resource pools](#)

[Configure roles](#)

[Configure accounts](#)

Dell PowerMax Storage admin role-based access control configuration

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

Resource pool configuration

Table 14: 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 Dell PowerMax Storage resources. 2. Select a specific SMTP server.
Jobs Tab	<ol style="list-style-type: none"> 1. Select the root level of all Dell PowerMax 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 Storage Utilization, select the Dell PowerMax Storage related reports.
Access Control Tab	<ol style="list-style-type: none"> 1. No Security Resources are assigned to this role.

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 the user must not have access to. 2. Name and submit the Resource Pool.
Role Configuration	<ol style="list-style-type: none"> 1. Select the permissions listed. 2. Name and submit the Role.
Dell PowerMax Storage Admin Account Configuration	<ol style="list-style-type: none"> 1. Create an account and link it to the newly created Dell PowerMax Storage Admin Resource Pool and Role.

Related information

[Configure resource pools](#)

[Configure roles](#)

[Configure accounts](#)

Configure tenants

A tenant is a grouping of resources and users that are administered by a tenant administrator. An IBM® Storage Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 [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 Defender 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 Defender 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 Defender 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 information

[Best practices for configuring tenants](#)

[Configure resource pools](#)

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 Defender 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 Defender 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 Defender 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 Defender 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 information

[Configure tenants](#)

[Configure resource pools](#)

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 Defender Copy Data Management. Credentials and keys are configured through the **Identities** view.

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

Related information

[Adding a key](#)

[Adding a credential](#)

Adding a key

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

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

IBM® Storage Defender 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 Defender 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.

Note: IBM® Storage Defender Copy Data Management does not support creating AWS key identities.

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 Defender Copy Data Management. Paste and save the generated public key to the `authorized_keys` file.
6. In IBM® Storage Defender 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 that is 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 Defender 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 Defender 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 Defender 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 Defender 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 following command:

```
$ ssh-keygen -t rsa -m PEM
```

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 Defender 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 Defender 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 Defender 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 Defender 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 Defender Copy Data Management appliance, the server to which the IBM® Storage Defender Copy Data Management appliance will connect needs to be registered. In this example, an Oracle server is used. In the IBM® Storage Defender 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 Defender 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 information

[Identities overview](#)

[Registering a storage provider](#)

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 information

[Identities overview](#)

[Registering a storage provider](#)

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 information

[Identities overview](#)

[Registering a storage provider](#)

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 Defender 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®, Dell PowerMax Storage, 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 [Registering a provider](#) and [Adding 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 Defender Copy Data Management, which include NetApp ONTAP storage controllers and clusters. See [Registering a provider](#).
- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Registering 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®, NetApp ONTAP, and Dell PowerMax Storage 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 information

[Creating an IBM Storage Virtualize Backup job definition](#)

[Creating a NetApp ONTAP Backup job definition](#)

[Creating a Pure Storage FlashArray Backup job definition](#)

[Creating 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. 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.
5. 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.

6. Click the source icon and select **Add VM Replication** from the context menu.
7. 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**.
8. 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 Defender Copy Data Management. The default initial label is VM Replication0.

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.

9. Enter a name for the new sub-policy (SLA Policy).
10. Click **Finish**. The policy gets created and listed under the **All SLA Policies** tab.

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 Defender 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.

5. Enter a name for the new sub-policy (SLA Policy).
6. Click **Finish**. The policy gets created and listed under the **All SLA Policies** tab.

Creating a Safeguarded Copy SLA policy

This procedure describes how to create a Safeguarded Copy (SGC) in IBM® Storage Defender 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 Defender 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 Defender 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**. The policy gets created and listed under the **All SLA Policies** tab.

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.

Before you begin

Important: The Global Mirror jobs are supported only on IBM FlashSystem version 8.7.0.3 or lower. Users cannot upgrade IBM FlashSystem to a higher version (versions greater than 8.7.0.3) until the Global Mirror configurations are migrated to policy-based replication (PBR). To understand the migration process, refer to the following links:

- [Converting Global Mirror legacy replication to policy-based asynchronous disaster recovery replication](#)
- [Policy-based replication is replacing Remote Copy \(Metro Mirror, Global Mirror, Global Mirror with Change Volumes and HyperSwap\)](#)

When you complete the migration and configure PBR/PBHA in IBM FlashSystem, complete the following steps in the Copy Data Management appliance:

1. Add an appropriate PBR/PBHA SLA policy. For more information, see [Creating a Policy-based replication \(PBR\) SLA policy](#) and [Creating a Policy-based high availability \(PBHA\) SLA policy](#).
2. Initiate backups by using the new PBR/PBHA SLA policy you have added.

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.

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

- b. Click **Add Global Mirror with Change Volumes**.
- c. 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**.
- d. 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 Defender 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 Defender 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.

FlashCopy 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.

- e. Enter a name for the new sub-policy (SLA Policy).
- f. Click **Finish**. The policy gets created and listed under the **All SLA Policies** tab.

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. 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 for Snapshot SLA policy.
5. 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.

6. 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.
7. 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 Defender 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.

8. Click **Finish**. The policy gets created and listed under the **All SLA Policies** tab.

Creating a Safeguarded Copy SLA policy

This procedure describes how to create a Safeguarded Copy (SGC) in IBM® Storage Defender 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 Defender 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 Defender 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**. The policy gets created and listed under the **All SLA Policies** tab.

Creating a policy-based high availability SLA policy

The procedure describes how to create a policy-based high availability (PBHA) SLA policy in IBM® Storage Defender Copy Data Management. Use this policy to create snapshots on both primary and secondary sites which are in a PBHA relationship. You can enable the PBHA three-site configuration to create snapshots on a disaster recovery (DR) site.

Before you begin

- A minimum IBM Storage FlashSystem® 9.1.0 version is required to use PBHA.
- You must configure PBHA in the source volumes to use the PBHA SLA policy in IBM® Storage Defender Copy Data Management.
- To use PBHA three-site configuration for backups, you must have a PBHA three-site setup in IBM Storage FlashSystem®.
- Enable the HA snapshots feature in the PBHA configuration. This ensures that the snapshots taken on the active management site are copied to the non-active management site.
- PBHA snapshots are application-consistent on the primary and target sites, and crash-consistent on DR sites.
- For PBHA storage volume backups, the consistency group option is disregarded. If the volumes belong to different volume groups, they are grouped by their respective volume groups. Individual snapshots are taken for each volume group.
- All create and delete operations are supported only on the Active Management Site (AMS) node. Therefore, volume revert from a snapshot is permitted exclusively on the current AMS node. To perform a revert operation, ensure that the backup selected corresponds to the node currently serving as the AMS node.
- According to IBM Storage FlashSystem® design, cloning a volume from a snapshot is not supported within a partition with PBHA replication policy enabled.
Due to this limitation, backups that are created by using the PBHA policy cannot be used to restore disks or databases on the original host. However, you can use the backups to replace the original volume with the contents of snapshot by selecting the **Revert** option in your restore job.
- For issues related to PBHA and PBR, see [Troubleshooting policy-based high availability and policy-based replication errors](#).

Important: If Global Mirror jobs exist and you want to migrate to PBHA configuration, refer to the following links:

- [Converting Global Mirror legacy replication to policy-based asynchronous disaster recovery replication](#)
- [Policy-based replication is replacing Remote Copy \(Metro Mirror, Global Mirror, Global Mirror with Change Volumes and HyperSwap\)](#)

Important: If you are using PBHA three-site SLA policies for your backups:

- All capabilities and limitations applicable to backups created by using a PBHA two-site SLA policy within a PBHA relationship also apply to backups created by using a PBHA three-site SLA policy.
- All capabilities and limitations applicable to backups created by using a PBR SLA policy also apply to backups on DR site created by using a PBHA three-site SLA policy.

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 **SLA Policies** pane, click **New**. The **New SLA Policy** pane opens.
4. Add an 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 5 minutes.

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

- b. Click **Add Policy Based HA (PBHA)**.
- c. In the **Options** pane, set the sub-policy options.

Keep Source and Target Snapshots

After some 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 (maximum)** field.

PBHA 3 site

Enable the **PBHA 3 site** option if you want to take additional snapshots on a DR site. To use the option, you must have a PBHA three-site setup in IBM® Storage Defender Copy Data Management.

Enter the age of the DR snapshot instances to purge in the **days** field, or the number of instances to keep in the **snapshots (maximum)** field.

Note:

The PBHA three-site configuration uses PBHA-based replication for primary and secondary sites with policy-based replication (PBR) for the DR site. In this setup, a single backup job performs the following actions:

- Backups are initiated on both sites by using PBHA, similar to the PBHA two-site configuration.
- Data is replicated to the DR site by using PBR.

Backups on PBHA sites are application-consistent, whereas the backups on the DR site by using PBR are crash-consistent.

- d. **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. 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 Defender Copy Data Management. The default name is PBHA0. **Take Safeguarded Copy** Enable the **Take Safeguarded Copy** option. When enabled, you can

see the **Perform Security Scan every** option. To enable the security scan, enable the option and select a security scan server from the list of available servers.

- e. Enter a name for a new sub-policy (SLA policy).
- f. Click **Finish**. A policy is created and listed under the **All SLA Policies** tab.

Creating a policy-based replication SLA policy

The procedure describes how to add a policy-based replication (PBR) policy to IBM® Storage Defender Copy Data Management. Use this policy to create snapshots on disaster recovery (DR) site, as per configuration in the replication policy.

Before you begin

- Volume group that contains the source volumes should have async DR policy associated with it before any backup job with PBR SLA policy is used to take backups.
- A minimum IBM Storage FlashSystem® 8.7.0 version is required to use PBR.
- The use of a single SLA policy with two sub-policies, one targeting the primary site and the other recovery site, is not supported in a PBR setup. If snapshots are required at both the primary and DR sites during the creation of a backup job, configure two distinct SLA policies, one for Snapshot or Safeguarded Copy and another for PBR.
- The IBM® Storage Defender Copy Data Management application does not attempt to create a replica on a remote storage array unless a replication relationship is already defined. You must configure the replication relationship from the storage UI before you can use it in IBM® Storage Defender Copy Data Management.
- In-place recovery, restore to source, and revert operations are not supported by PBR at the VM and application levels.
- Any remote array snapshot for PBR is only crash-consistent. Snapshots on the primary site are application-consistent.
- If a consistency group is enabled for PBR backup of storage volumes, all volumes must be in the same volume group.
- For backup of application (database) volumes, all selected databases must be in the same volume group.
- For issues related to PBHA and PBR, see [Troubleshooting policy-based high availability and policy-based replication errors](#).

Important: If Global Mirror jobs exist and you want to migrate to PBR configuration, refer to the following links to understand the migration process:

- [Converting Global Mirror legacy replication to policy-based asynchronous disaster recovery replication](#)
- [Policy-based replication is replacing Remote Copy \(Metro Mirror, Global Mirror, Global Mirror with Change Volumes and HyperSwap\)](#)

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 Policy** 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 Policy Based Replications (PBR)**.
- c. In the **Options** pane, set the sub-policy options.

Keep Target Snapshots

After some 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 (maximum)** 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 Defender Copy Data Management. The default name is PBR0.

Take Safeguarded Copy

Enable the **Take Safeguarded Copy** option.

- d. Enter a name for the new sub-policy (SLA Policy).
- e. Click **Finish**. The policy gets created and listed under the **All SLA Policies** tab.

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 Defender 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 Replication**.
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 Defender Copy Data Management. The default initial label is VM Replication0.

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 Defender 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 Defender 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 Defender 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 Defender 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 Replication**.
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 Defender Copy Data Management. The default initial label is VM Replication0.

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 Defender Copy Data Management. The default initial name is Snapshot0.

3. If your Pure Storage FlashArray supports CloudSnap functionality, you can add a snapshot offload, which creates 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 Defender Copy Data Management user interface indicates that a Pure Storage FlashArray CloudSnap offload job has been completed even though the transfer is still occurring in the background, which depends on network speed. Consider setting an age as the retention for offload copies when you use the Pure Storage FlashArray CloudSnap functionality to ensure that a sufficient amount of time has passed for data to be transferred to the S3 storage target or NFS share before it must be condensed out from a backup. This is important if several offload jobs are run in quick succession.

Offload Target

Select the offload target.

Note: IBM® Storage Defender Copy Data Management does not support the **Cloud** option as a snapshot 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 Defender 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 Defender 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.

Configure Dell PowerMax Storage SLA policies

The procedures describe how to configure Dell PowerMax Storage SLA policies.

Adding a snapshot sub-policy to a Dell PowerMax Storage SLA policy

The procedure describes steps to add a snapshot sub-policy to a Dell PowerMax Storage policy.

1. Click the **Configure** tab. On the **Views** pane, select **Sites & Providers**, then expand the Dell PowerMax Storage and add the supported storage.
2. On the **Views** pane, select **SLA Policies**. The **All SLA Policies** pane opens.
3. In the **SLA Policies** pane, click **New**. The **New SLA Policies** pane opens.
4. The set frequency and interval of an SLA policy.
Select the source icon and set the recovery point objective (RPO) to define the minimum backup frequency and interval. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, and then set the interval in the **Interval** field. The lowest available frequency is 5 minutes.

Important: Any changes to the frequency and interval of an SLA policy apply to all associated job schedules.

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

Keep Snapshots

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

Snapshot Prefix Label

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

Note: The snapshot name character limit for Dell PowerMax Storage is 32 characters. If the snapshot name exceeds the limit, the snapshot procedure fails.

Name

Enter a name to replace the default snapshot sub-policy name, which is displayed in IBM® Storage Defender Copy Data Management. The default initial name is Snapshot0.

6. Click **Finish** to create the SLA policy.

Tip: You can use the add snapshot sub-policy to create a snapshot on the R1 leg of Metro 2-site configuration.

Adding a remote replication sub-policy to a Dell PowerMax Storage SLA policy

Use this procedure to add a remote replication sub-policy to a Dell PowerMax Storage SLA Policy.

Before you begin

- Register source and destination Dell PowerMax Storage systems in IBM® Storage Defender Copy Data Management. For more information, see [Registering a Dell PowerMax Storage provider](#).
- Add the Symmetrix Remote Data Facility (SRDF) group to the source volume storage group. To create an SRDF group, see [SRDF groups](#).

Procedure

1. Click the **Configure** tab. On the **Views** pane, select **Sites & Providers**, then expand the Dell PowerMax Storage and add the supported storage.
2. On the **Views** pane, select **SLA Policies**. The **All SLA Policies** pane opens.
3. In the **SLA Policies** pane, click **New**. The **New SLA Policies** pane opens.
4. The set frequency and interval of an SLA policy.
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 5 minutes.

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

5. Click **Add Remote Replication**.
6. In the **Replication Mode** pane, select a replication mode from the list of modes such as Synchronous, Asynchronous, or Metro.

Tip: Select the **Metro Replication Mode** to create a snapshot on the R2 leg of the Metro 2-site configuration.

7. In the **Replication Destination** pane, select a Dell PowerMax Storage host destination from the list of available resources as the destination. If the Metro replication mode is selected, confirm that the **Replication Destination** corresponds to the R2 leg of the Metro 2-site configuration.

Tip: To create snapshots on the R1 leg of a Metro 2-site configuration, add a snapshot sub-policy when you create an SLA. For more information, see [Adding a snapshot sub-policy to a Dell PowerMax SLA policy](#).

8. In the **Options** pane, set the replication sub-policy options.
Keep Snapshots

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

Snapshot Prefix Label

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

Note: The snapshot name character limit for Dell PowerMax Storage is 32 characters. If the snapshot name exceeds the limit, the snapshot procedure fails.

Name

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

9. Review and verify the SLA policy information, and click **Finish** to save the SLA policy. The SLA policy is displayed in the **All SLA Policies** pane. You can apply the SLA policy to backup job definitions.

Configure Dell PowerFlex Storage SLA policies

The procedures describe how to configure Dell PowerFlex Storage SLA policies.

Adding a snapshot sub-policy to a Dell PowerFlex Storage SLA policy

The procedure describes steps to add a snapshot sub-policy to a Dell PowerFlex Storage policy.

1. Click the **Configure** tab. On the **Views** pane, select **Sites & Providers**, then expand the Dell PowerFlex Storage and add the supported storage.
2. On the **Views** pane, select **SLA Policies**. The **All SLA Policies** pane opens.
3. In the **SLA Policies** pane, click **New**. The **New SLA Policies** pane opens.
4. The set frequency and interval of an SLA policy.
Select the source icon and set the recovery point objective (RPO) to define the minimum backup frequency and interval. In the **Frequency** field select Minutes, Hourly, Daily, Weekly, or Monthly, and then set the interval in the **Interval** field. The lowest available frequency is 5 minutes.

Important: Any changes to the frequency and interval of an SLA policy apply to all associated job schedules.

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

Keep Snapshots

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

Snapshot Prefix Label

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

Note: The snapshot name character limit for Dell PowerFlex Storage is 32 characters. If the snapshot name exceeds the limit, the snapshot procedure fails.

Name

Enter a name to replace the default snapshot sub-policy name, which is displayed in IBM® Storage Defender Copy Data Management. The default initial name is Snapshot0.

6. Click **Finish** to create the SLA policy.

Adding a remote replication sub-policy to a Dell PowerFlex Storage SLA policy

Use this procedure to add a remote replication sub-policy to a Dell PowerFlex Storage SLA Policy.

Before you begin

- Register source and destination Dell PowerFlex Storage systems in IBM® Storage Defender Copy Data Management. For more information, see [Registering a Dell PowerFlex Storage provider](#).
- Add the Remote Consistency Group (RCG) group to the source volume storage group. To create a RCG group, see [RCG groups](#).

Procedure

1. Click the **Configure** tab. On the **Views** pane, select **Sites & Providers**, then expand the Dell PowerFlex Storage and add the supported storage.
2. On the **Views** pane, select **SLA Policies**. The **All SLA Policies** pane opens.
3. In the **SLA Policies** pane, click **New**. The **New SLA Policies** pane opens.
4. The set frequency and interval of an SLA policy.
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 5 minutes.

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

5. Click **Add Remote Replication**.
6. In the **Replication Mode**, the **Asynchronous** option is the only available mode and cannot be edited.
7. In the **Replication Destination** pane, select a Dell PowerFlex Storage host destination from the list of available resources as the destination.
8. In the **Options** pane, set the replication sub-policy options.

Keep Snapshots

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

Snapshot Prefix Label

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

Note: The snapshot name character limit for Dell PowerFlex Storage is 32 characters. If the snapshot name exceeds the limit, the snapshot procedure fails.

Name

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

9. Review and verify the SLA policy information, and click **Finish** to save the SLA policy. The SLA policy is displayed in the **All SLA Policies** pane. You can apply the SLA policy to backup job definitions.

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 by 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.

What to do next

As the next step, you need to specify the **Parameters**, **Identify**, and **Application Server** fields.

- Click the **Parameters** field. Add a parameter to the script, then click **Add**. If the script does not require any parameters, leave the field empty.
- Click the **Identity** field to add or create the credentials that you need to run the script. The field is mandatory when you define a script.
- Click the **Application Server** field to define the server where you need to run the script. The field is mandatory when you define a script.

Note: To add additional parameters to a script, enter a parameter in the field, then click **Add**. For parameter examples, see [“Using state and status arguments in postscripts” on page 252](#).

Related information

[Job definition overview](#)

[Using state and status arguments in postscripts](#)

[Return code reference](#)

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 information

[Job definition overview](#)

[Using state and status arguments in postscripts](#)

[Return code reference](#)

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 information

[Job definition overview](#)

[Using state and status arguments in postscripts](#)

[Return code reference](#)

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 information

[Editing a schedule](#)

[Deleting a schedule](#)

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 information

[Creating a schedule](#)

[Deleting a schedule](#)

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 information

[Creating a schedule](#)

[Editing a schedule](#)

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 Defender 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 Defender 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 information

[Job definition overview](#)

[Monitor a job session](#)

[Start, pause, and hold a job session](#)

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 Defender 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 information

[Creating a 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 information

[Creating a 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 information

[Creating a 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 information

[Creating a 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 information

[Start, pause, and hold a job session](#)

[Collecting logs for troubleshooting](#)

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 information

[Start, pause, and hold a job session](#)

[Collecting logs for troubleshooting](#)

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 information

[Start, pause, and hold a job session](#)

[Collecting logs for troubleshooting](#)

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 Defender 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 Defender 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 Defender Copy Data Management appliance. Use the script job to add functionality to IBM® Storage Defender 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 Defender 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 Defender Copy Data Management catalog, and removes related snapshots.

Related information

[Configure SLA policies](#)

[Creating a schedule](#)

[Monitor a job session](#)

[Creating a Database Inventory job definition](#)

[Creating a File System Inventory job definition](#)

[Creating an IBM Storage Virtualize Inventory job definition](#)

[Creating a NetApp ONTAP Storage Inventory job definition](#)

[Creating a NetApp ONTAP File Inventory job definition](#)

[Creating a Pure Storage FlashArray Inventory job definition](#)

[Creating a VMware Inventory job definition](#)
[Creating an InterSystems Database Backup job definition](#)
[Creating an SAP HANA Backup job definition](#)
[Creating an Oracle Backup job definition](#)
[Creating a SQL Backup job definition](#)
[Creating a File System Backup job](#)
[Creating an IBM Storage Virtualize Backup job definition](#)
[Creating a NetApp ONTAP Backup job definition](#)
[Creating a Pure Storage FlashArray Backup job definition](#)
[Creating a VMware Backup job definition](#)
[Creating an InterSystems Database Restore job definition](#)
[Creating an SAP HANA Restore job definition](#)
[Creating an Oracle Restore job definition](#)
[Creating a Microsoft SQL Restore job definition](#)
[Creating a File System Restore job definition](#)
[Creating an IBM Storage Virtualize Restore job definition](#)
[Creating a NetAPP ONTAP Restore job definition](#)
[Creating a Pure Storage FlashArray Restore job definition](#)
[Creating a VMware 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 Defender 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**

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 information

[Editing a job definition](#)
[Deleting a job definition](#)
[Creating a 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 Defender 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click Add to add it to the list.

Click **Ok**.

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 information

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a 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 Defender 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 Defender 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 information

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a 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 Defender 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 Defender 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

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 Defender 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

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 information

[Creating a NetApp ONTAP File Inventory job definition](#)

[Editing 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

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 information

[Creating a NetApp ONTAP Storage Inventory job definition](#)

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

[Search and filter guidelines](#)

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 Defender 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

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 information

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a 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 Defender 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

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 information

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

Creating a Dell PowerMax Storage Inventory job definition

A Dell PowerMax Storage Inventory job provides the framework to catalog and collect information about high-level Dell PowerMax Storage objects on your Dell PowerMax Storage systems.

Before you begin

- At least one application provider must be associated with a database inventory job definition, such as an Oracle resource. Define an inventory job, and then add application providers. See [Register a Provider](#).
- For email notifications, at least one SMTP server must be configured. Define a job and then 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 that are 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 Defender Copy Data Management.

A Dell PowerMax Storage Inventory job provides the framework to catalog and collect information about high-level Dell objects on your Dell storage systems. You can select one or more providers in a single job definition for cataloging.

Run a Dell PowerMax Storage Inventory job to analyze your Dell environment in real time and navigate and correlate the objects from across the Enterprise in a single view. Also, you can infer sprawl, storage overutilization, and other storage inefficiencies.

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, and then select **Dell PowerMax**.
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 a job definition using default options, click **Create Job**. You can run the job manually from the **Jobs** tab.
6. To edit options before a job definition creation, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that you can do on a provider.

Number of catalog instances to keep

Older high-level Dell objects are automatically removed from the inventory after a job is ran a certain number of times. Specify how many times a job must be run to retain high-level Dell object before older objects are purged.

7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select an SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add the email addresses to the list.

Click **Ok**.

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 notification about the status of each task is sent when the job is complete.

Related information

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

Creating a Dell PowerFlex Storage Inventory job definition

A Dell PowerFlex Storage Inventory job provides the framework to catalog and collect information about high-level Dell PowerFlex Storage objects like system info, volumes and hosts on your registered Dell PowerFlex Storage systems.

Before you begin

- At least one storage system must be registered as a provider. Define an inventory job and add storage providers. See [Registering a Dell PowerFlex Storage provider](#).
- Configure at least one SMTP server for email notifications. Define a job, and then add SMTP resources. See [Registering an SMTP provider](#).

Considerations:

- 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](#).

About this task

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

A Dell PowerFlex Storage Inventory job provides the framework to catalog and collect information about high-level Dell PowerFlex Storage systems. You can select one or more providers in a single job definition for cataloging.

Run a Dell PowerFlex Storage Inventory job to analyze your Dell PowerFlex Storage environment in real time and navigate and correlate the objects from across the enterprise in a single view. Also, you can infer sprawl, storage overutilization, and other storage inefficiencies.

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, and then select **Dell PowerFlex**.
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 a job definition using default options, click **Create Job**. You can run the job manually from the **Jobs** tab.
6. To edit options before a job definition creation, click **Advanced**. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent cataloging tasks that you can do on a provider.

Number of catalog instances to keep

Older high-level Dell PowerFlex Storage objects are automatically removed from the inventory after a job runs a certain number of times. Specify how many times a job must be run to retain high-level Dell PowerFlex Storage objects before older objects are purged.

7. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select an SMTP Server to use for job status email notifications. If an SMTP server is not selected, an email is not sent.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add the email addresses to the list.

Click **Ok**.

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 information

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a 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 procedure describes how to create an InterSystems Database Backup job definition. You can create InterSystems Caché and InterSystems IRIS application servers in IBM® Storage Defender 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 Defender Copy Data Management provides application database copy management through application-consistent backup creation, cloning, and recovery. IBM® Storage Defender 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 Defender 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. Select only **one SLA** policy from the list of available SLA policies.
 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.

To execute scripts on a remote Linux server, see [Custom selection of a directory for script execution on a remote Linux server](#).

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

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 information

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

[Creating an InterSystems Database Restore job definition](#)

Creating an InterSystem IRIS application (on AIX®) scanning backup job

The procedure describes 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 Defender 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 Defender 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 Defender 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 Defender Copy Data Management are displayed.
7. Choose the appropriate proxy host(s) for IBM® Storage Defender 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. Select only **one SLA** policy from the list of available SLA policies.
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](#).

For more information, see [Running the post snapshot script before the security scan enhancement](#).

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.

To execute scripts on a remote Linux server, see [Custom selection of a directory for script execution on a remote Linux server](#).

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

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 an SAP HANA Backup job definition

This procedure describes how to create an SAP HANA Backup job definition.

Before you begin

- Register a storage provider before you create a backup job. See [Registering an Application Server - SAP HANA](#).
 - 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 Defender Copy Data Management provides application database copy management through application-consistent backup creation, and recovery. IBM® Storage Defender 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 Defender 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 Defender 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.

While backing up an SAP HANA System Replication (SAP HSR) Cluster, selecting the database, a snapshot backup is completed using the existing Primary system only.

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. Select only **one SLA** policy from the list of available SLA policies.
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.

To execute scripts on a remote Linux server, see [Custom selection of a directory for script execution on a remote Linux server](#).

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

10. To edit Log Backup options before creating the job definition, click Log Backup. If Backup Logs is selected, IBM® Storage Defender 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 an SAP HANA Restore job definition. See [Create a Restore Job Definition - SAP HANA](#).

Related information

[Editing a job definition](#)
[Deleting a job definition](#)
[Creating a schedule](#)
[Creating an SAP HANA Restore job definition](#)

Creating an Oracle Backup job definition

The procedure describes 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 [Creating a Database Inventory job definition - Oracle](#).
 - Configure an SLA Policy. See [SLA Policies](#).
 - Review Oracle requirements. See [Oracle requirements](#).
- Oracle database considerations:**
- To ensure that filesystem permissions are retained correctly when IBM® Storage Defender Copy Data Management moves Oracle data between servers, ensure that the user and group IDs of the Oracle users (e.g. oracle, oinstall, dba) are consistent across all the servers. Refer to Oracle documentation for recommended uid and gid values.
 - If Oracle data resides on LVM volumes, you must stop and disable the lvm2-lvmetad service before running Backup or Restore jobs. Leaving the service enabled can prevent volume groups from being resignatured correctly during restore and can lead to data corruption if the original volume group is also present on the same system. To disable the lvm2-lvmetad service, run the following commands:

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

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

```
use_lvmetad = 0
```

- 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 Defender 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 Defender Copy Data Management automatically purges the logs under this new location after each successful backup. For IBM® SVC, IBM® Storage Defender 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 Defender Copy Data Management purges the logs after

the FlashCopy® operation only. For databases that are protected only by a Global Mirror workflow, IBM® Storage Defender 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 Defender 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.
- Configure at least one SMTP server for email notifications. Define a job, and then add SMTP resources. See [Registering an SMTP provider](#).
- 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](#).

About this task

IBM® Storage Defender Copy Data Management provides application database copy management through application-consistent backup creation, cloning, and recovery. IBM® Storage Defender 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 Defender 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 Defender 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 Defender 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. Select only **one SLA** policy from the list of available SLA policies.
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 Defender 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](#).

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.

To execute scripts on a remote Linux server, see [Custom selection of a directory for script execution on a remote Linux server](#).

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

10. To edit Log Backup options before creating the job definition, click **Log Backup**. If **Create additional archive log destination** is selected, IBM® Storage Defender 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 Defender Copy Data Management automatically discovers the location where Oracle writes archived logs. If this location resides on storage from a supported vendor, IBM® Storage Defender 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 Defender 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 Defender 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 Defender 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 Defender Copy Data Management-managed destination. Even in this case, IBM® Storage Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender Copy Data Management Restore jobs.

Tip: User-defined masking scripts that were in existence before IBM® Storage Defender 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 necessary, start the job session immediately rather than waiting for the scheduled job completion. See [Start, Pause, and Hold a Job Session](#).
- Click **Jobs** tab to monitor the progress of the job session. See [Monitor a Job Session](#).
- Enable notification options, to send an email about the status of each task when the job completes.
- Use the **Inventory Browse** to review the recovery point. See [Browse Inventory](#).

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 Defender 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 Defender 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 Defender 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 Defender Copy Data Management can utilize the IBM® Storage Defender 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 Defender Copy Data Management are displayed.
7. Choose the appropriate proxy host(s) for IBM® Storage Defender 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. Select only **one SLA** policy from the list of available SLA policies.
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 Defender 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.

To execute scripts on a remote Linux server, see [Custom selection of a directory for script execution on a remote Linux server](#).

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

12. To edit Log Backup options before creating the job definition, click **Log Backup**. If **Create additional archive log destination** is selected, IBM® Storage Defender 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 Defender Copy Data Management automatically discovers the location where Oracle writes archived logs. If this location resides on storage from a supported vendor, IBM® Storage Defender 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 Defender 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 Defender 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 Defender 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 Defender Copy Data Management-managed destination. Even in this case, IBM® Storage Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender Copy Data Management Restore jobs.

Tip: User-defined masking scripts that were in existence before IBM® Storage Defender 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 procedure describes 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 [Creating a Database Inventory job definition](#).
- Configure an SLA Policy. See [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.
- 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 Defender 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 Defender 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 preceeding Backup job ran.
 - Configure at least one SMTP server for email notifications. Define a job, and then add SMTP resources. See [Registering an SMTP provider](#).
 - 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.

About this task

IBM® Storage Defender Copy Data Management provides application database copy management through application-consistent backup creation, cloning, and recovery. IBM® Storage Defender 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 Defender 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 Defender 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. Select only **one SLA** policy from the list of available SLA policies.
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.

To execute scripts on a remote Linux server, see [Custom selection of a directory for script execution on a remote Linux server](#).

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

11. To edit Log Backup options before creating the job definition, click **Log Backup**. If **Backup Logs** is selected, IBM® Storage Defender 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 Defender 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 Defender Copy Data Management automatically truncates post log backups of databases that it backs up. If database logs are not backed up with IBM® Storage Defender Copy Data Management, logs are not truncated by IBM® Storage Defender 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 necessary, start the job session immediately rather than waiting for the scheduled job completion. See [Start, Pause, and Hold a Job Session](#).
- Click **Jobs** tab to monitor the progress of the job session. See [Monitor a Job Session](#).
- Enable notification options, to send an email about the status of each task when the job completes.
- Use the **Inventory Browse** to review the recovery point. See [Browse Inventory](#).
- Create a SQL Application Restore job definition. See [Create a Restore Job Definition - SQL](#).

Creating a File System Backup job

The procedure describes 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 [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 preceeding Backup job ran.
 - Configure at least one SMTP server for email notifications. Define a job, and then add SMTP resources. See [Registering an SMTP provider](#).
 - 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](#).

About this task

IBM® Storage Defender Copy Data Management provides application database copy management of physical Windows™, Linux®, and AIX® file systems through a File System Backup job.

Note: Supported filesystem types:

- AIX - JFS2
- Windows - NTFS
- Linux - xfs, ext3, and ext4

Procedure

1. Click the **Jobs** tab. Expand the **File System** folder, and 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. Select only **one SLA** policy from the list of available SLA policies.
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.

Enable Quiesce

Enable this option to quiesce the file system during the backup process. By default, the option is disabled. When enabled, the backup job will quiesce the file system before creating a snapshot (pre-snapshot operation) to ensure a consistent backup copy. After the snapshot is taken, the backup job will unquiesce the file system during the post-snapshot operation.

Note: The /root and /tmp file systems on Linux® and AIX® and the C : drive on Windows™, are excluded from both quiesce and unquiesce operations.

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.

To execute scripts on a remote Linux server, see [Custom selection of a directory for script execution on a remote Linux server](#).

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

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 necessary, start the job session immediately rather than waiting for the scheduled job completion. See [Start, Pause, and Hold a Job Session](#).
- Click **Jobs** tab to monitor the progress of the job session. See [Monitor a Job Session](#).
- Enable notification options, to send an email about the status of each task when the job completes.
- Use the **Inventory Browse** to review the recovery point. See [Browse Inventory](#).
- Create an OS Volume Restore job definition. See [Create a Restore Job Definition - File System](#).

Creating an IBM® Storage Virtualize Backup job definition

The procedure describes 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](#).
- Configure at least one SMTP server for email notifications. Define a job, and then add SMTP resources. See [Registering an SMTP provider](#).
- IBM® providers utilize port 22 for communication with IBM® Storage Defender 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 that are defined in the schedule. See [Create a Schedule](#).

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. Select only **one SLA** policy from the list of available SLA policies.
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.

To execute scripts on a remote Linux server, see [Custom selection of a directory for script execution on a remote Linux server](#).

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

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 necessary, start the job session immediately rather than waiting for the scheduled job completion. See [Start, Pause, and Hold a Job Session](#).
- Click **Jobs** tab to monitor the progress of the job session. See [Monitor a Job Session](#).
- Enable notification options, to send an email about the status of each task when the job completes.
- Use the **Inventory Browse** 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**.
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.

To execute scripts on a remote Linux server, see [Custom selection of a directory for script execution on a remote Linux server](#).

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

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 procedure describes 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](#).
- Configure at least one SMTP server for email notifications. Define a job, and then add SMTP resources. See [Registering an SMTP provider](#).
- IBM® providers utilize port 22 for communication with IBM® Storage Defender 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 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.

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. Select only **one SLA** policy from the list of available SLA policies.
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.

To execute scripts on a remote Linux server, see [Custom selection of a directory for script execution on a remote Linux server](#).

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

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 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](#).
- Register a storage provider before you create a backup job. See [Registering an IBM Storage Virtualize for snapshot provider](#).
- Configure at least one SMTP server for email notifications. Define a job, and then add SMTP resources. See [Registering an SMTP provider](#).

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.

To execute scripts on a remote Linux server, see [Custom selection of a directory for script execution on a remote Linux server](#).

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

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 necessary, start the job session immediately rather than waiting for the scheduled job completion. See [Start, Pause, and Hold a Job Session](#).
- Click **Jobs** tab to monitor the progress of the job session. See [Monitor a Job Session](#).
- Enable notification options, to send an email about the status of each task when the job completes.
- Use the **Inventory Browse** to review the recovery point. See [Browse Inventory](#).

Creating a policy-based high availability backup

This procedure describes how to create a policy-based high availability (PBHA) backup.

Before you begin

- Register a storage provider before you create a backup job. See [Registering an IBM Storage Virtualize for snapshot provider](#).
- Create a PBHA policy that you want to use for PBHA backup jobs. For more information, see [Creating a policy-based high availability SLA policy](#).

- Backups created using the PBHA policy cannot be used to perform disk or database restores on the original host.
- PBHA snapshots can be used to revert the original instance or restore to an alternate host, provided it is of the same type and version.
- Restoring a SQL database from a remote copy is not supported when the backup is created using a PBHA policy. Only local disk-based restores are permitted. Since PBHA snapshot cannot be restored to partition having PBHA policy, it cannot be used to restore on the original host. Restoring on an alternate host makes the snapshot a remote copy, which is not supported. Hence, only option available is to revert the original snapshot when performing the database restore. Disk restore to an alternate host is supported.
- Restoring to a different SQL server that is connected to the DR site storage is also unsupported due to an existing limitation of SQL Server.
- Restoring a database to an alternate SAP HANA server is not supported. Since PBHA snapshot cannot be restored to partition having PBHA policy, it cannot be used to restore on the original host. Restoring on alternate host makes the snapshot a remote copy which is not supported. Hence only option available is to revert original snapshot when performing database restore.
- For PBHA storage volume backups, the consistency group option is disregarded. If the volumes belong to different volume groups, they are grouped by their respective volume groups. Individual snapshots are taken for each volume group.
- Do not run multiple separate backup jobs on a given application at the same time.
- For issues related to PBHA and PBR, see [Troubleshooting policy-based high availability and policy-based replication errors](#).

Important: If you are using PBHA three-site SLA policies for your backups:

- All capabilities and limitations applicable to backups created by using a PBHA two-site SLA policy within a PBHA relationship also apply to backups created by using a PBHA three-site SLA policy.
- All capabilities and limitations applicable to backups created by using a PBR SLA policy also apply to backups on DR site created by using a PBHA three-site SLA policy.

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.

Note: This option is not applicable for PBHA policy-based backups. In such backups, volumes are grouped into volume groups, and a separate snapshot is created for each 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.

To execute scripts on a remote Linux server, see [Custom selection of a directory for script execution on a remote Linux server](#).

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.

- 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

- Review and verify the job specific information job-specific information, and click **Create Job**. The job runs according to the specified schedule. You can also run the job manually from the **Jobs** tab.

What to do next

- If necessary, start the job session immediately rather than waiting for the scheduled job completion. See [Start, Pause, and Hold a Job Session](#).
- Click **Jobs** tab to monitor the progress of the job session. See [Monitor a Job Session](#).
- Enable notification options, to send an email about the status of each task when the job completes.
- Use the **Inventory Browse** to review the recovery point. See [Browse Inventory](#).

Creating a policy-based replication backup

The procedure describes how to create a policy-based replication (PBR) backup.

Before you begin

- Register a storage provider before you create a backup job. See [Registering an IBM Storage Virtualize for Snapshot provider](#).
- Create a PBHA policy that you want to use for PBHA backup jobs. For more information, see [Creating a policy-based high availability SLA policy](#).
- Application backups taken by using PBR snapshots are crash consistent only.
- Restoring a SQL database from a remote copy is not supported when the backup is created using a PBR policy. Only local disk-based restores are permitted. Even if the database resides on a host with access to both primary and disaster recovery (DR) site storage, the snapshot is taken at the DR site, rendering it a remote copy from the host's perspective.
- Restoring to a different SQL server that is connected to the DR site storage is not supported due to an existing limitation of SQL Server.
- Restoring a database to an alternate SAP HANA server is not supported. Therefore, when a backup is created by using a PBR policy, you can only use it to restore the database on the original server that is connected to the DR storage instance. Restoring the backup to a different database server, even if it is connected to the DR storage, is not supported.
- If a consistency group is enabled for PBR backup of storage volumes, all volumes must be in the same volume group.
- For backup of application (database) volumes, all selected databases must be in the same volume group.
- Do not run multiple separate backup jobs on a specific application at a time.
- For issues related to PBHA and PBR, see [Troubleshooting policy-based high availability and policy-based replication errors](#).

Procedure

- Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **IBM® Storage Virtualize for Snapshot**.
- Click **New**, then select **Backup**. The job editor opens.
- 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 applicable SLA policy of the data backup.
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.

To execute scripts on a remote Linux server, see [Custom selection of a directory for script execution on a remote Linux server](#).

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 an SMTP server to enable job status email notifications. If an SMTP server is not selected, an email is not sent.

Recipients

Enter the email addresses of the recipients for email notifications. Click **Add** to add it to the list.

Click **OK**.

10. Review and verify the job specific information job-specific information, and click **Create Job**. The job runs according to the specified schedule. You can also run the job manually from the **Jobs** tab.

What to do next

- If necessary, start the job session immediately rather than waiting for the scheduled job completion. See [Start, Pause, and Hold a Job Session](#).
- Click **Jobs** tab to monitor the progress of the job session. See [Monitor a Job Session](#).
- Enable notification options, to send an email about the status of each task when the job completes.
- Use the **Inventory Browse** to review the recovery point. See [Browse Inventory](#).

Creating a NetApp ONTAP Backup job definition

The procedure describes 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 Defender 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. Select only **one SLA** policy from the list of available SLA policies.
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.

To execute scripts on a remote Linux server, see [Custom selection of a directory for script execution on a remote Linux server](#).

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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 information

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

[Creating an NetAPP ONTAP Restore job definition](#)

Creating a Pure Storage FlashArray Backup job definition

The procedure describes 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 Defender 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. Select only **one SLA** policy from the list of available SLA policies.
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.

To execute scripts on a remote Linux server, see [Custom selection of a directory for script execution on a remote Linux server](#).

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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 information

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

[Creating a Pure Storage FlashArray Restore job definition](#)

Creating a VMware Backup job definition

The procedure describes 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 Defender Copy Data Management was tested against VMware Tools 12.3.5.
 - 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 Defender 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 Defender 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 FlashCopy and FlashCopy NoCopy backup jobs on the same source volume.
- The Sentinel scanning capability requires iSCSI or Fibre Channel (FC) connectivity between the IBM Storage FlashSystem® and the IBM® Storage Defender Sentinel host.
 - When using Fibre Channel, make sure the Sentinel host is zoned and defined in the IBM Storage FlashSystem® array as Fibre Channel host.
 - When using iSCSI, make sure the Sentinel host has network reachability over iSCSI port (3260/TCP) with IBM Storage FlashSystem® and the Sentinel host is defined as iSCSI host.

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, Dell PowerMax Storage, 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. Select only **one SLA** policy from the list of available SLA policies.

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 Defender Copy Data Management generates logs pertaining to the application log truncation function and copies them to the following location on the IBM® Storage Defender 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.

To execute scripts on a remote Linux server, see [Custom selection of a directory for script execution on a remote Linux server](#).

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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 information

[Editing a job definition](#)

[Deleting a job definition](#)

[Configure scripts](#)

[Creating a schedule](#)

[Creating VMware Backup job proxies](#)

[Creating a VMware Restore job definition](#)

Creating a Dell PowerMax Storage backup job definition

The procedure describes how to create a Dell PowerMax Storage backup job definition.

Before you begin

- Create and run a Dell PowerMax Storage inventory job that includes the providers you wish to back up. See [Creating a Dell PowerMax Storage Inventory job definition](#).
- Configure an SLA policy. See [Configure SLA Policies](#).
- For email notifications, at least one SMTP server must be configured. Define a job and then 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 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.

About this task

Back up Dell data with a snapshot by using a Dell PowerMax Storage Backup job. The RPO and copy data parameters are defined in an SLA Policy, which is then applied to the Backup job definition, along with a specified activation time to meet your copy data criteria. Supported sources include Dell PowerMax Storage systems.

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **Dell PowerMax**.
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. Select only **one SLA** policy from the list of available SLA policies.
7. Click **Create Job**, to create a job definition by using default options. You can run the job manually from the **Jobs** tab.

8. Click **Advanced** to edit options before a job definition creation. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent transfers between a source and a destination.

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select an SMTP server to enable job status email notifications. If an SMTP server is not selected, an email is not sent.

Recipients

Enter the email addresses of the recipients of status email notifications. Click **Add** to add it to the list.

Click **OK**.

10. Review and verify the job specific information job-specific information, and click **Create Job**. The job runs according to the specified schedule. You can also run the job manually from the **Jobs** tab.

What to do next

- If necessary, start the job session immediately rather than waiting for the scheduled job completion. See [Start, Pause, and Hold a Job Session](#).
- Click **Jobs** tab to monitor the progress of the job session. See [Monitor a Job Session](#).
- Enable notification options, to send an email about the status of each task when the job completes.
- Use the **Inventory Browse** to review the recovery point. See [Browse Inventory](#).
- Create a Dell PowerMax Storage restore job definition. See [Creating a Restore job definition - Dell PowerMax Storage](#).

Related information

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

[Creating a restore job definition - Dell PowerMax Storage](#)

Creating a Dell PowerFlex Storage backup job definition

The procedure describes how to create a Dell PowerFlex Storage backup job definition.

Before you begin

- Create and run a Dell PowerFlex Storage inventory job that includes the providers you wish to back up. See [Creating a Dell PowerFlex Storage Inventory job definition](#).
- Configure an SLA policy. See [Configure SLA Policies](#).
- For email notifications, at least one SMTP server must be configured. Define a job and then 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 that are defined in the schedule. See [Create a Schedule](#).

About this task

Back up Dell data with a snapshot by using a Dell PowerFlex Storage Backup job. The RPO and copy data parameters are defined in an SLA Policy, which is then applied to the Backup job definition, along with a specified activation time to meet your copy data criteria. Supported sources include Dell PowerFlex Storage systems.

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **Dell PowerFlex**.

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. Select only **one SLA** policy from the list of available SLA policies.
7. Click **Create Job**, to create a job definition by using default options. You can run the job manually from the **Jobs** tab.
8. Click **Advanced** to edit options before a job definition creation. Set the job definition options.

Maximum concurrent tasks

Set the maximum number of concurrent transfers between a source and a 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.

9. Optionally, expand the **Notification** section to select the job notification options.

SMTP Server

From the list of available SMTP resources, select an SMTP server to enable job status email notifications. If an SMTP server is not selected, an email is not sent.

Recipients

Enter the email addresses of the recipients of status email notifications. Click **Add** to add it to the list.

Click **OK**.

10. Review and verify the job specific information job-specific information, and click **Create Job**. The job runs according to the specified schedule. You can also run the job manually from the **Jobs** tab.

What to do next

- If necessary, start the job session immediately rather than waiting for the scheduled job completion. See [Start, Pause, and Hold a Job Session](#).
- Click **Jobs** tab to monitor the progress of the job session.. See [Monitor a Job Session](#).
- Enable notification options, to send an email about the status of each task when the job completes.
- Use the **Inventory Browse** to review the recovery point. See [Browse Inventory](#).
- Create a Dell PowerFlex Storage restore job definition. See [Creating a Restore job definition - Dell PowerFlex Storage](#).

Related information

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

[Creating a restore job definition - Dell PowerFlex Storage](#)

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 procedure describes how to create an InterSystems Database Restore job definition. You can create InterSystems Caché and InterSystems IRIS application servers in IBM® Storage Defender 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 Defender 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 Defender 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.

Note: Remote snapshot restores require the ESXi host to have access to both the local and remote PowerFlex arrays. If hosts at each site are isolated by cluster and cannot access both arrays, the restore workflow may fail or behave unexpectedly. Select an application backed by the remote array as the restore target in such scenarios.

About this task

IBM® Storage Defender 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 Defender Copy Data Management Instant Disk Restore jobs. IBM® Storage Defender 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 Disk Restore**.

Note: Instant Database Restore is not supported for InterSystems Database due to limitation of the database itself.

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.

Note: User can select recovery points as the local or remote copy as applicable for relevant storages supporting recovery from remote snapshots. Revert (restore to original production volume) from local and remote snapshot is not supported for InterSystems Databases on any storage systems.

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. For **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](#).

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 Defender 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

Disabled - Never make permanent

User Selection - Allows the user to select Make Permanent or Cleanup when the job session is pending

Note: In the case of Dell PowerFlex Storage, the clone copy created after the **Make Permanent** operation remains the Snapshot type, not the Volume type. If you want to delete the original source volume from Dell PowerFlex Storage, ensure that you do not delete the Copy Data Management created and mapped clone copies that have been made permanent.

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.

Note: This option is not applicable in the case of Dell PowerFlex Storage and can be ignored.

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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.

Important: For Dell PowerFlex Storage based virtual InterSystems restores:

- Ensure that the VMware ESXi hostnames registered in Dell PowerFlex Storage exactly match the hostnames configured on the VMware vSphere environment. If you see a mismatch in the names, rename the ESXi host Dell PowerFlex Storage to match with the names on the vCenter.
- Ensure that the Dell PowerFlex Storage SDC component is installed and running on ESXi hosts where the application VMs are hosted.
- If your VM is hosted on an ESXi host which belongs to an ESXi host cluster, ensure that the peer ESXi hosts in the cluster are also registered with the same name on Dell PowerFlex Storage.

- 13.

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 information

[Creating an InterSystems Database Backup job definition](#)

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

Creating an SAP HANA Restore job definition

This procedure describes how to create restore jobs for SAP HANA.

Before you begin

- Create and run an 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 an 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.

Note: Remote snapshot restores require the ESXi host to have access to both the local and remote PowerFlex arrays. If hosts at each site are isolated by cluster and cannot access both arrays, the restore workflow may fail or behave unexpectedly. Select an application backed by the remote array as the restore target in such scenarios.

About this task

IBM® Storage Defender 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 Defender Copy Data Management Instant Disk Restore jobs. IBM® Storage Defender 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.

Note: User can select recovery points as the local or remote copy as applicable for relevant storages supporting recovery from remote snapshots.
Revert (restore to original production volume) from remote snapshot is not supported for SAP HANA database on any storage systems.

Note: For SAP HANA restore jobs, the SAP HANA database host is not available for restore if the host is removed from the inventory job, even though backup copies of that SAP HANA database host are available.
User should re-add the removed SAP HANA host database to the inventory job and run it again to see backup copies available for the restore job.

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.

Note: The Restore (Instant Disk and Instant Database) operations for SAP HANA are not supported on an alternate host. You need to restore it on the same source host.

To restore a database on an SAP HANA System Replication (SAP HSR) cluster, complete the following steps for each restore option:

- **Instant database restore:** Suspend SAP HANA system replication before initiating the restore to prevent automatic failover. Stop the Pacemaker service to ensure that no failover is triggered during the operation.
- **Instant disk restore:** Suspend SAP HANA system replication before initiating the restore to prevent automatic failover. Stop the Pacemaker service to ensure that no failover is triggered during the operation.

For more information, refer to the [SAP HANA HSR Cluster Server Support FAQ](#) section.

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
 - 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 Defender 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: The **Revert** function is available only for snapshots created by using the IBM® Storage Virtualize, IBM Storage Virtualize for Snapshot, Dell PowerMax Storage, and Dell PowerFlex Storage SLA policies.

Note: In the case of Dell PowerFlex Storage, the clone copy created after the **Make Permanent** operation remains the Snapshot type, not the Volume type. If you want to delete the original source volume from Dell PowerFlex Storage, ensure that you do not delete the Copy Data Management created and mapped clone copies that have been made permanent.

- 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.

Note: This option is not applicable in the case of Dell PowerFlex Storage and can be ignored.

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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.

Important: For Dell PowerFlex Storage based virtual SAP HANA restores:

- Ensure that the VMware ESXi hostnames registered in Dell PowerFlex Storage exactly match the hostnames configured on the VMware vSphere environment. If you see a mismatch in the names, rename the ESXi host Dell PowerFlex Storage to match with the names on the vCenter.
- Ensure that the Dell PowerFlex Storage SDC component is installed and running on ESXi hosts where the application VMs are hosted.
- If your VM is hosted on an ESXi host which belongs to an ESXi host cluster, ensure that the peer ESXi hosts in the cluster are also registered with the same name on Dell PowerFlex Storage.

- 13.

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 information

[Creating a File System Backup job](#)

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

Creating an Oracle Restore job definition

The procedure describes how to create an restore jobs for Oracle.

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 Defender Copy Data Management moves Oracle data between servers, ensure that the user and group IDs of the Oracle users (e.g. oracle, oinstall, dba) are consistent across all the servers. Refer to Oracle documentation for recommended uid and gid values.
- If Oracle data resides on LVM volumes, you must stop and disable the lvm2-lvmetad service before running Backup or Restore jobs. Leaving the service enabled can prevent volume groups from being resignatured correctly during restore and can lead to data corruption if the original volume group is also present on the same system. To disable the lvm2-lvmetad service, run the following commands:

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

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

```
use_lvmetad = 0
```

- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Register a Provider](#).
- 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.”

Note: Remote snapshot restores require the ESXi host to have access to both the local and remote PowerFlex arrays. If hosts at each site are isolated by cluster and cannot access both arrays, the restore workflow may fail or behave unexpectedly. Select an application backed by the remote array as the restore target in such scenarios.

About this task

IBM® Storage Defender 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 Defender Copy Data Management Instant Disk Restore jobs. IBM® Storage Defender 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.

Note: User can select recovery points as the local or remote copy as applicable for relevant storages supporting recovery from remote snapshots. Revert (restore to original production volume) from remote snapshot is not supported for Oracle database on any storage systems.

6. Click **Copy**. Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. If you would like to choose a specific version or use latest successful scan, select a site then click **Select Specific Version or Use Latest Successful Scan**. 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 Defender 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

Put database in read only mode after restore

This option is available for Instant Database Restore workflows.

Select this option to put the database in read only mode after restore.

Put database in noarchivelog mode after restore

This option is available for Instant Database Restore workflows.

Select this option to put the database in noarchivelog mode after restore.

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 Defender 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 Defender 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.

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 Defender 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 Defender 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.

Prepare scripts post Instant Disk Restore

Select this option to create scripts at the end of the Instant Disk Restore job for Oracle. These scripts can be used by database administrators to perform Oracle database restore. After the mount operations completes, IBM® Storage Defender Copy Data Management creates RMAN scripts and SQL scripts on the Copy Data Management. The scripts are created in the /data/log/ecxdeployer/<YYYY-MM-DD>/<random_string> folder. To recover the database, users must execute these scripts.

The script files are named as Step-**<number>**_<operation>.<extension>, where:

- <number> refers the sequence in which the file must be run.
- <operation> refers to the specific task that the script performs.
- <extension> refers the type of file, either RMAN or SQL.

Depending on the type of script, execute the script in RMAN or by using SQLPlus. The following script files are generated:

- `Step-1_DatabaseRename.sql`: The file contains the control file creation script that verifies the paths of the data file and log files are correct.
- `Step-2_RecoverDatabase.rman`: The file contains the steps to recover the database by using the RMAN script.
- `Step-3_RecoverDatabase.sql`: The file contains the steps to recover the database by using the SQL script. If the restore operation is performed by using the RMAN script file (`Step-2_RecoverDatabase.rman`), the step is not required.
- `Step-4_OpenDatabase.sql`: The file contains the script that opens the database post the restore operation.
- `Step-5_createTempTableSpace.sql`: The file contains the script to verify the temp file path.
- `Step-6_dropTempTableSpace.sql`: The file contains the script (optional script) that drops the temporary tablespace created by using the `Step-5_createTempTableSpace.sql` script.

For more information about the running the script files, see [Oracle requirements](#).

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 Defender 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 Defender 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 old session

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

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 Defender 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

Disabled - Never make permanent

User Selection - Allows the user to select Make Permanent or Cleanup when the job session is pending

Revert

Set the source production Oracle machines to the snapshot. The option can have an impact on the database downtime and presents a risk to the source production machine.

Enabled - Always revert

Disabled - Never revert

Note:

The **Revert** function is available only for snapshots created by using the IBM® Storage Virtualize, IBM Storage Virtualize for Snapshot, Dell PowerMax Storage, and Dell PowerFlex Storage SLA policies.

The **Revert** function is not supported for AIX based Oracle deployments.

The **Revert** function is not supported for LVM based Oracle deployments.

Attention: Do not choose the **Revert** option when multiple databases are hosted on the same disks (volumes or LUNs) or data stores. For more information on supported configurations, see the [Oracle requirements](#).

Note: In the case of Dell PowerFlex Storage, the clone copy created after the **Make Permanent** operation remains the Snapshot type, not the Volume type. If you want to delete the original source volume from Dell PowerFlex Storage, ensure that you do not delete the Copy Data Management created and mapped clone copies that have been made permanent.

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.

Note: This option is not applicable in the case of Dell PowerFlex Storage and can be ignored.

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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.

Important: For Dell PowerFlex Storage based virtual Oracle restores:

- Ensure that the VMware ESXi hostnames registered in Dell PowerFlex Storage exactly match the hostnames configured on the VMware vSphere environment. If you see a mismatch in the names, rename the ESXi host Dell PowerFlex Storage to match with the names on the vCenter.
- Ensure that the Dell PowerFlex Storage SDC component is installed and running on ESXi hosts where the application VMs are hosted.
- If your VM is hosted on an ESXi host which belongs to an ESXi host cluster, ensure that the peer ESXi hosts in the cluster are also registered with the same name on Dell PowerFlex Storage.

13.

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 information

[Creating an Oracle Backup job definition](#)

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

Creating a Microsoft™ SQL Restore job definition

The procedure describes how to create restore jobs for Microsoft™ SQL.

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 Defender Copy Data Management leverages a Windows™ proxy server to temporarily mount disks from snapshots, generate a new signature, then mount to original server.
 - 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.
- Install Microsoft Multipath IO (MPIO) role by using the Add Roles and Features wizard in Windows Server Manager on the SQL virtual machine.

Note: Remote snapshot restores require the ESXi host to have access to both the local and remote PowerFlex arrays. If hosts at each site are isolated by cluster and cannot access both arrays, the restore workflow may fail or behave unexpectedly. Select an application backed by the remote array as the restore target in such scenarios.

About this task

IBM® Storage Defender 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 Defender Copy Data Management Instant Disk Restore jobs. IBM® Storage Defender 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 Defender 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 Defender 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. For Microsoft SQL (Standalone and Failover Cluster), available template options are **Instant Database Restore** and **System DB Restore**. For Microsoft SQL (Always On Availability Group), available template options are **Instant Disk Restore**, **Instant Seeding**, and **Instant Database Restore**.

Note: User can select recovery points as the local or remote copy as applicable for relevant storages supporting recovery from remote snapshots. Revert (restore to original production volume) from remote snapshot is not supported for SQL database on any storage systems.

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.

Note: Microsoft™ SQL Standalone, SQL Fail-over, and SQL Availability Group Database restores from the secondary and remote arrays are not supported.

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™ proxy server in the **Select a Windows server to resignature LUNs** section. When resignaturing a copy, IBM® Storage Defender 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 FCI cluster. It is recommended to either select a standalone virtual or physical Windows™ node as a proxy server or another Standalone SQL server 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 Defender 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 Defender 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 Defender 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 Defender 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 old session

Enabling this option allows a scheduled session of a recovery job to force an existing pending session to clean up associated resources so the new session can run. Disable this option to keep an existing test environment running without being cleaned up.

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 Defender 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

Disabled - Never make permanent

User Selection - Allows the user to select Make Permanent or Cleanup when the job session is pending

Revert

Set the source production SQL machines to the snapshot. This may affect database availability and pose a potential risk to the source production machine.

Enabled - Always revert

Disabled - Never revert

Note: The **Revert** function is available only for snapshots created by using the IBM® Storage Virtualize, IBM Storage Virtualize for Snapshot, Dell PowerMax Storage, and Dell PowerFlex Storage SLA policies.

Attention: Do not choose the **Revert** option where multiple databases are hosted on the common disks (volumes/LUNs) or data stores. For more information on supported configuration refer to the Revert SQL Machine section in [Microsoft SQL requirements](#).

Note: In the case of Dell PowerFlex Storage, the clone copy created after the **Make Permanent** operation remains the Snapshot type, not the Volume type. If you want to delete the original source volume from Dell PowerFlex Storage, ensure that you do not delete the Copy Data Management created and mapped clone copies that have been made permanent.

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.

Note: This option is not applicable in the case of Dell PowerFlex Storage and can be ignored.

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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.

Important: For Dell PowerFlex Storage based virtual Microsoft SQL restores:

- Ensure that the VMware ESXi hostnames registered in Dell PowerFlex Storage exactly match the hostnames configured on the VMware vSphere environment. If you see a mismatch in the names, rename the ESXi host Dell PowerFlex Storage to match with the names on the vCenter.
- Ensure that the Dell PowerFlex Storage SDC component is installed and running on ESXi hosts where the application VMs are hosted.
- If your VM is hosted on an ESXi host which belongs to an ESXi host cluster, ensure that the peer ESXi hosts in the cluster are also registered with the same name on Dell PowerFlex Storage.

14.

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 information

[Creating a SQL Backup job definition](#)

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

Creating a File System Restore job definition

The procedure describes how to create restore jobs for File System.

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 Defender 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 Defender Copy Data Management Instant Disk Restore jobs. IBM®

Storage Defender 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.

Note: User can select recovery points as the local or remote copy as applicable for relevant storages supporting recovery from remote snapshots.
Revert (restore to original production volume) from remote snapshot is not supported for File Systems on any storage systems.

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 Defender 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

Disabled - Never make permanent

User Selection - Allows the user to select Make Permanent or Cleanup when the job session is pending

Revert

Set the source production File System to the snapshot. The option can have an impact on the database downtime and presents a risk to the source production machine.

Enabled – Always revert

Enabled – Never revert

User Selection – Allows the user to select Revert or Cleanup when the job session is completed

Note: The **Revert** function is only available for backups created by using the IBM Storage Virtualize, IBM Storage Virtualize for Snapshot, and Dell PowerMax Storage SLA policies.

Attention: The **Revert** function is supported only for Linux and Windows file systems.

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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 information

[Creating a File System Backup job](#)

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

Creating an IBM® Storage Virtualize Restore job definition

The procedure describes how to create restore jobs for IBM® Storage Virtualize.

IBM® Storage Defender 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 Defender 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 Defender 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 procedure describes 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](#).
- Configure at least one SMTP server for email notifications. Define a job, and then add SMTP resources. See [Registering an SMTP provider](#).
- IBM® providers utilize port 22 for communication with IBM® Storage Defender 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 that are 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.

- a. 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 session

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.

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

10. 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.

11. 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 a Restore Volume(s) IBM® Storage Virtualize Restore job definition

The procedure describes how to create a Restore Volume(s) IBM® Storage Virtualize Restore job definition.

Before you begin

- Configure at least one SMTP server for email notifications. Define a job, and then add SMTP resources. See [Registering an SMTP provider](#).
- IBM® providers utilize port 22 for communication with IBM® Storage Defender 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 that are 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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 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](#).
- Configure at least one SMTP server for email notifications. Define a job, and then add SMTP resources. See [Registering an SMTP provider](#).

- **Warning:** The Safeguarded copies will be removed from storage after the retention period has expired as defined in storage SLA. IBM® Storage Defender 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.

- IBM® providers utilize port 22 for communication with IBM® Storage Defender 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. 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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

The procedure describes how to create restore jobs for an IBM® Storage Virtualize for Snapshot.

IBM® Storage Defender 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 Defender 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 procedure describes 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 Snapshot](#).
- Configure at least one SMTP server for email notifications. Define a job, and then add SMTP resources. See [Registering an SMTP provider](#).
- IBM® providers utilize port 22 for communication with IBM® Storage Defender 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 that are 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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 necessary, start the job session immediately rather than waiting for the scheduled job completion. See [Start, Pause, and Hold a Job Session](#).
- Click **Jobs** tab to monitor the progress of the job session. See [Monitor a Job Session](#).

- Enable notification options, to send an email about the status of each task when the job completes.
- Use the **Inventory Browse** to review the recovery point. See [Browse Inventory](#).

Creating a Restore Volume(s) IBM® Storage Virtualize for Snapshot Restore job definition

The procedure describes how to create a Restore Volume(s) IBM® Storage Virtualize for Snapshot Restore job definition. The workflow allows user to restore/revert the production volume to the snapshot.

Before you begin

- For email notifications, at least one SMTP server must be configured. Before defining a job, add SMTP resources. See [Registering an SMTP provider](#).
Considerations:
- IBM® providers utilize port 22 for communication with IBM® Storage Defender Copy Data Management.
- 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 **Restore Volume(s)** template.
5. Click **Source**. From the drop-down menu select **Volume** to select a source site and an associated IBM® source to view volumes with available recovery points. Select one or more resources, and change the order in which the resources are recovered by dragging and dropping the resources in the grid. Alternatively, select **Volume Search** from the drop-down menu to search for volumes with available recovery points. Add volume copies to the job definition by clicking **Add**. Change the order in which the resources are recovered by dragging and dropping the resources in the grid.
6. Click **Copy**. Sites containing copies of the selected data display. Select a site. By default the latest copy of your data is used. To choose a specific version, select a site and click **Select Version**. Click the **Version field** to view specific copies and their associated job and completion time. If recovery from one snapshot fails, another copy from the same site is used.

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**.
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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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 information

[Creating an IBM Storage Virtualize for Snapshot Backup job definition](#)

[Using state and status arguments in postscripts](#)

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

[Search and filter guidelines](#)

Creating an NetAPP ONTAP Restore job definition

The procedure describes how to create restore jobs for NetAPP ONTAP.

IBM® Storage Defender 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 Defender 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 Defender 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 Defender 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 procedure describes 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 Defender 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 **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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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 information

[Creating a NetApp ONTAP Backup job definition](#)

[Using state and status arguments in postscripts](#)

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

Creating a Restore Volume(s) NetApp ONTAP Restore job definition

The procedure describes 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 Defender 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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 information

[Creating a NetApp ONTAP Backup job definition](#)

[Using state and status arguments in postscripts](#)

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

[Search and filter guidelines](#)

Creating a Restore File(s) NetApp ONTAP Restore job definition

The procedure describes 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 Defender 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list. Click **OK**.

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 information

[Creating a NetApp ONTAP Backup job definition](#)

[Using state and status arguments in postscripts](#)

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

[Search and filter guidelines](#)

Creating a Pure Storage FlashArray Restore job definition

The procedure describes how to create restore jobs for Pure Storage FlashArray.

IBM® Storage Defender 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 Defender 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 Defender 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 Defender 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 procedure describes 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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 information

[Creating a Pure Storage FlashArray Backup job definition](#)

[Using state and status arguments in postscripts](#)

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

[Search and filter guidelines](#)

Creating a Restore Volume(s) Pure Storage FlashArray Restore job definition

The procedure describes 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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 information

[Creating a Pure Storage FlashArray Backup job definition](#)

[Using state and status arguments in postscripts](#)

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

[Search and filter guidelines](#)

Creating a VMware Restore job definition

The procedure describes how to create restore jobs for VMware.

IBM® Storage Defender 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 Defender 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.

Important: For Dell PowerFlex Storage based VM restores:

- Ensure that the VMware ESXi hostnames registered in Dell PowerFlex Storage exactly match the hostnames configured on the VMware vSphere environment. If you see a mismatch in the names, rename the ESXi host Dell PowerFlex Storage to match with the names on the vCenter.
- Ensure that the Dell PowerFlex Storage SDC component is installed and running on ESXi hosts where the application VMs are hosted.
- If your VM is hosted on an ESXi host which belongs to an ESXi host cluster, ensure that the peer ESXi hosts in the cluster are also registered with the same name on Dell PowerFlex Storage.

Creating an Instant Disk Restore VMware Restore job definition

The procedure describes 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 Defender 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 Defender 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 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 Defender 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 Defender Copy Data Management does not require the VASA provider metadata for Restore workflows, as IBM® Storage Defender 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.

Note: When restoring VM disks to a remote Dell PowerFlex Storage array, ensure that the ESXi hosts at Site A and Site B have access to both the local and remote arrays. If the hosts reside in separate clusters and do not share connectivity to both arrays, the Instant Disk Restore workflow requires the following additional configuration:

- **For Instant Disk Restore:** Uncheck the default **Use original VM** option and select a VM that is hosted on the remote ESXi cluster as the restore target.

If this is not configured correctly, the restore job may fail or produce unexpected results.

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.

Datastores

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.

Note: This option is not applicable in the case of Dell PowerFlex Storage and can be ignored.

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.

Note: In the case of Dell PowerFlex Storage, the clone copy created after the **Make Permanent** operation remains the Snapshot type, not the Volume type. If you want to delete the original source volume from Dell PowerFlex Storage, ensure that you do not delete the Copy Data Management created and mapped clone copies that have been made permanent.

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 a to 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.

Revert

The **Revert** option allows administrators to restore VMware datastore volumes by using snapshots created during backup operations. When enabled, the option reverts the production volume on the datastore to a selected snapshot that is maintained by Copy Data Management. The selected snapshot is mounted, and its point-in-time state replaces the current contents of the volume.

Enabled - Always revert

Disabled - Never revert

Note: The **Revert** function is used to restore a datastore to a previous state and is only supported for datastores, not for virtual machines (VMs). This functionality is supported on Dell PowerMax Storage, Dell PowerFlex Storage, and IBM Storage Virtualize for Snapshotstorage controllers.

Attention: Verify that all dependent VMs are powered off to prevent data corruption. The operation is destructive and overwrites the current contents of the datastore.

The **Revert** function is available only for datastores and is not supported for individual VMs. To use this function, confirm that the source snapshot is created by using a Copy Data Management restore job.

Before initiating a restore job with the **Revert** function enabled, perform the following checks:

- Verify the restore point to ensure it is valid and appropriate for the operation.
- Power off all dependent VMs to prevent potential data corruption during the revert process.

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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 information

[Creating a VMware Backup job definition](#)

[Using state and status arguments in postscripts](#)

https://library.netapp.com/ecm/ecm_download_file/ECMP1197114

https://library.netapp.com/ecm/ecm_download_file/ECMP1287548

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

[Search and filter guidelines](#)

Creating an Instant VM Restore VMware Restore job definition

The procedure describes 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 Defender Copy Data Management was tested against VMware Tools 12.3.5.
- 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 Defender 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 Defender 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 Defender Copy Data Management does not require the VASA provider metadata for Restore workflows, as IBM® Storage Defender 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.

Note: When restoring VM to a remote Dell PowerFlex Storage array, ensure that the ESXi hosts at Site A and Site B have access to both the local and remote arrays. If the hosts reside in separate clusters and do not share connectivity to both arrays, the Instant VM Restore workflow requires the following additional configuration:

- **For Instant VM Restore:** Select the **Use alternative host or cluster** option.

If this is not configured correctly, the restore job may fail or produce unexpected results.

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 Defender 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 Defender Copy Data Management appliance and the server, and installs an IBM® Storage Defender 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 Defender 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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 information

[Creating a VMware Backup job definition](#)

[Using state and status arguments in postscripts](#)

https://library.netapp.com/ecm/ecm_download_file/ECMP1197114

https://library.netapp.com/ecm/ecm_download_file/ECMP1287548

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

[Search and filter guidelines](#)

Creating an Instant VM Restore (Long Distance) VMware Restore job definition

The procedure describes 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 Defender Copy Data Management was tested against VMware Tools 12.3.5.
- 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 Defender 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 Defender 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 Defender Copy Data Management does not require the VASA provider metadata for Restore workflows, as IBM® Storage Defender 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 Defender Copy Data Management appliance and the server, and installs an IBM® Storage Defender 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 Defender 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.

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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **OK**.

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 information

[Creating a VMware Backup job definition](#)

[Using state and status arguments in postscripts](#)

https://library.netapp.com/ecm/ecm_download_file/ECMP1197114

https://library.netapp.com/ecm/ecm_download_file/ECMP1287548

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

[Search and filter guidelines](#)

Creating a Dell PowerMax Storage restore job definition

The procedure describes how to create restore jobs for Dell PowerMax Storage.

IBM® Storage Defender Copy Data Management leverages Copy Data Management technology for recovering Dell PowerMax Storage volumes through Dell PowerMax Storage Restore jobs.

Instant Disk Restore

Provides the instant writable access to a volume. An IBM® Storage Defender Copy Data Management snapshot is mapped to a target server where it can be accessed, copied, or used in production immediately, if required.

Restore Volume(s)

Recovers a volume from a snapshot that is created through a Dell PowerMax Storage Backup job. You can restore the volumes to the original location.

Creating an Instant Disk Restore Dell PowerMax Restore job definition

Create an Instant Disk Restore Dell PowerMax Storage restore job definition to restore backups created by using the Dell PowerMax Storage jobs.

Before you begin

- Create and run one or more Dell PowerMax Storage backup jobs. See [Creating a Dell PowerMax Storage Backup job definition](#).
- Configure at least one SMTP server for email notifications. Define a job, and then add SMTP resources. See [Registering an SMTP provider](#).
Considerations:
 - 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](#).

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **Dell PowerMax**.
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 list, select **Volume** to select a source site and an associated Dell PowerMax 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 list 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.

Note: In Dell PowerMax Storage array-based restore jobs, you can select the local or remote copy recovery points.

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 Dell PowerMax Storage destination.
8. Click **Create Job**, to create a job definition by using default options. You can run the job manually from the **Jobs** tab.
9. Click **Advanced**, to edit options before a job definition creation. Set the job definition options.

Continue with next source on failure

If enabled, the restore job continues with the next resource in case the recovery of a resource 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 cleanup of pending old sessions

Enable this option to allow a scheduled recovery job to force cleanup of resources from a pending session, so a new session can run. Disable this option to keep an existing test environment running without automatic cleanup.

Job-Level Scripts

You can run job-level pre-scripts and post-scripts 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. 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**. 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 is injected and run.

Repeat the procedure to add more 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

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. Review and verify the job specific information job-specific information, and click **Create Job**. The job runs according to the specified schedule. You can also run the job 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 information

[Creating a Dell PowerMax Storage backup job definition](#)

[Using state and status arguments in postscripts](#)

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

[Search and filter guidelines](#)

Creating a Restore Volume(s) Dell PowerMax Storage Restore job definition

The procedures describe how to create a Restore Volume(s) Dell PowerMax Storage Restore job definition.

Before you begin

- Create and run one or more Dell PowerMax Storage backup jobs. See [Creating a Dell PowerMax Storage backup job definition](#).
- Configure at least one SMTP server for email notifications. See [Registering an SMTP provider](#).
- Add SMTP resources and define a job. “[Registering a storage provider](#)” on page 40.

Note: 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](#).

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **Dell PowerMax**.
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 Dell PowerMax 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.

Note: In Dell PowerMax Storage array based restore jobs, the user can select recovery points as the local or remote copy.
IBM® Storage Defender Copy Data Management does not support in-place restore from a remote copy in case of Dell PowerMax Storage array based restores.

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**. The **Restore to original volume** option is selected by default and cannot be edited.
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.

- a. 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. Ensure that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure scripts](#).
- b. 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**.

- c. Click the **Identity** field to add or create the credentials that are required to run the script.
 - d. Click the **Application Server** field to define the location where the script is injected and run.
 - e. Run the script again to add more Pre-Scripts and Post-Scripts. For information about script return codes, see [Return code reference](#).
 - f. 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.
 - g. Select **Continue operation on script failure** to continue running the restore job if a command in any of the scripts that are associated with the job failure.
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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

11. Optionally, expand the **Schedule** section to select one or more 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. When each schedule is selected, the schedule name and description is displayed.

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. Review and verify the job specific information job-specific information, and click **Create Job**. The job runs according to the specified schedule. You can also run the job manually from the **Jobs** tab.

What to do next

- If you do not want to wait until the next scheduled job runs, 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 notification with information about the status of each task is sent when the job completes.

Related information

[Creating a Dell PowerMax Storage backup job definition](#)

[Using state and status arguments in postscripts](#)

[Editing a job definition](#)

[Deleting a job definition](#)

Creating a Dell PowerFlex Storage restore job definition

The procedure describes how to create restore jobs for Dell PowerFlex Storage.

IBM® Storage Defender Copy Data Management leverages Copy Data Management technology for recovering Dell PowerFlex Storage volumes through Dell PowerFlex Storage Restore jobs.

Instant Disk Restore

Provides the instant writable access to a volume. An IBM® Storage Defender Copy Data Management creates a clone from the snapshot and is mapped to a target host server where it can be accessed, copied, or used in production immediately, if required.

Restore Volume(s)

Recovers a volume from a snapshot that is created through a Dell PowerFlex Storage Backup job. You can restore the volumes to the original location. This operation is also known as Revert.

Creating an Instant Disk Restore Dell PowerFlex Storage Restore job definition

Create an Instant Disk Restore Dell PowerFlex Storage restore job definition to restore backups created by using the Dell PowerFlex Storage jobs.

Before you begin

- Create and run one or more Dell PowerFlex Storage backup jobs. See [Creating a Dell PowerFlex Storage Backup job definition](#).
 - Configure at least one SMTP server for email notifications. Define a job, and then add SMTP resources. See [Registering an SMTP provider](#).
- Considerations:**
- 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](#).

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **Dell PowerFlex**.
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 list, select **Volume** to select a source site and an associated Dell PowerFlex 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 list 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.

Note: In Dell PowerFlex Storage array-based restore jobs, you can select the local or remote copy recovery points.

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 Dell PowerFlex Storage destination.

Note: If you chose a remote copy in the previous step, ensure the destination is the same remote Dell PowerFlex Storage array. Restoring from a remote destination to a local is not supported.

8. Click **Create Job**, to create a job definition by using default options. You can run the job manually from the **Jobs** tab.
9. Click **Advanced**, to edit options before a job definition creation. Set the job definition options.

Make IA® clone resource permanent

Enable to convert the snapshot copy into a permanent resource that is not removed when the Instant Disk Restore job completes.

Continue with next source on failure

If enabled, the restore job continues with the next resource in case the recovery of a resource 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 cleanup of pending old sessions

Enable this option to allow a scheduled recovery job to force cleanup of resources from a pending session, so a new session can run. Disable this option to keep an existing test environment running without automatic cleanup.

Job-level scripts

You can run job-level pre-scripts and post-scripts 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. 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**. 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 is injected and run.

Repeat the procedure to add more 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

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. Review and verify the job specific information job-specific information, and click **Create Job**. The job runs according to the specified schedule. You can also run the job manually from the **Jobs** tab.
13. If the **Make IA® clone resource permanent** option is disabled, the Disk Restore job completes with the **Resource Active** state. In this case, right-click the job session listed in the **Activity** tab to view relevant options, such as **IR (make permanent)**, **End IA volume (Cleanup)** and **End IA volume (Without Cleanup)**. If you select the **IR (make permanent)** option, the cloned copy becomes a standalone snapshot that the Copy Data Management appliance no longer monitors for retention.

Note: In the case of Dell PowerFlex Storage, the clone copy created after the **Make Permanent** operation remains the Snapshot type, not the Volume type. If you want to delete the original source volume from Dell PowerFlex Storage, ensure that you do not delete the Copy Data Management created and mapped clone copies that have been made permanent.

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 information

[Creating a Dell PowerFlex Storage backup job definition](#)

[Using state and status arguments in postscripts](#)

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

[Search and filter guidelines](#)

Creating a Restore Volume(s) Dell PowerFlex Storage Restore job definition

The procedures describe how to create a Restore Volume(s) Dell PowerFlex Storage Restore job definition.

Before you begin

- Create and run one or more Dell PowerFlex Storage backup jobs. See [Creating a Dell PowerFlex Storage backup job definition](#).
- Configure at least one SMTP server for email notifications. See [Registering an SMTP provider](#).
- Add SMTP resources and define a job. “[Registering a storage provider](#)” on page 40.

Note: 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](#).

Note: When performing a restore job for a volume that is part of a Remote Consistency Group (RCG) in a Dell PowerFlex Storage, you need to disable RCG on the Dell PowerFlex Storage array. To disable RCG, see [PowerFlex Replication How to Unpair from RCG](#) and [Terminate replication in an RCG](#).

Procedure

1. Click the **Jobs** tab. Expand the **Storage Controller** folder, then select **Dell PowerFlex**.
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 Dell PowerFlex 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.

Note: In Dell PowerFlex Storage array based restore jobs, the user can select recovery points as the local or remote copy.
IBM® Storage Defender Copy Data Management does not support in-place restore from a remote copy in case of Dell PowerFlex Storage array based restores.

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**. The **Restore to original volume** option is selected by default and cannot be edited.
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.

- a. 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. Ensure that scripts can also be uploaded and edited through the **Scripts** view on the **Configure** tab. See [Configure scripts](#).
- b. 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**.

- c. Click the **Identity** field to add or create the credentials that are required to run the script.
- d. Click the **Application Server** field to define the location where the script is injected and run.
- e. Run the script again to add more Pre-Scripts and Post-Scripts. For information about script return codes, see [Return code reference](#).

- f. 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.
 - g. Select **Continue operation on script failure** to continue running the restore job if a command in any of the scripts that are associated with the job failure.
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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.
11. Optionally, expand the **Schedule** section to select one or more 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. When each schedule is selected, the schedule name and description is displayed.

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. Review and verify the job specific information job-specific information, and click **Create Job**. The job runs according to the specified schedule. You can also run the job manually from the **Jobs** tab.

What to do next

- If you do not want to wait until the next scheduled job runs, 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 notification with information about the status of each task is sent when the job completes.

Related information

[Creating a Dell PowerFlex Storage backup job definition](#)

[Using state and status arguments in postscripts](#)

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

[Search and filter guidelines](#)

Renaming mount points and initialization parameter options

The topic describes how to rename mount points and initialize 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 Defender Copy Data Management appends a timestamp to the original mount point. For example:

Table 15: Mount Point Rename - Append a timestamp	
Original Mount Point	New Mount Point
/u02/myproddb/data	/u02/myproddb/data_1479400505
+MYPRODDATA	+MYPRODDATA1479400505

Do not rename

Select this option if you do not want to rename mount points or ASM diskgroups during recovery. IBM® Storage Defender 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:

Table 16: Mount Point Rename - Add a custom prefix		
Original Mount Point	Prefix	New Mount Point
/u02/myproddb/data	tmp	/tmpu02/myproddb/data
/u02/myproddb/data	/tmp/	/tmp/u02/myproddb/data
+MYPRODDATA	/tmp/	+TMPMYPRODDATA

Add a custom suffix

Select this option and specify a custom suffix to be appended the source paths/names. For example:

Table 17: Mount Point Rename - Add a custom suffix		
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 Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender Copy Data Management depending on the appropriate values for the destination.

Related information

[Creating an Oracle Restore job definition](#)

[Creating an InterSystems Database Restore job definition](#)

Using state and status arguments in postscripts

The topic describes how to use 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 Defender 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 information

[Creating an IBM Storage Virtualize Restore job definition](#)

[Creating an NetAPP ONTAP Restore job definition](#)

[Creating a VMware Restore job definition](#)

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 Defender Copy Data Management when a job is deleted. The cleanup procedure reclaims space on your storage devices, cleans up your IBM® Storage Defender 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 is deleted. All logs associated with the deleted job are removed from IBM® Storage Defender 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,

After deleting a pending application, IBM®, Dell PowerMax Storage, Dell PowerFlex Storage, NetApp ONTAP, Pure Storage FlashArray, or VMware Backup or Restore 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®, Dell PowerMax Storage, Dell PowerFlex Storage, NetApp ONTAP, Pure Storage FlashArray, or VMware Backup or Restore job, all associated IBM®, Dell PowerMax Storage, Dell PowerFlex Storage, NetApp ONTAP, Pure Storage FlashArray, and VMware locations are removed by the Maintenance job. Once the Maintenance job completes, IBM®, Dell PowerMax Storage, Dell PowerFlex Storage, NetApp ONTAP, Pure Storage FlashArray, and 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®, Dell PowerMax Storage, Dell PowerFlex Storage, 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 Defender Copy Data Management interface.

Related information

[Deleting 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.

Recipients

Enter the email addresses of the status email notifications recipients. Click **Add** to add it to the list.

Click **Ok**.

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 information

[Report Overview](#)

[Editing a job definition](#)

[Deleting a job definition](#)

[Creating a schedule](#)

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 information

[Deleting 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 information

[Editing a job definition](#)

[Maintenance job](#)

Searching

The topics in the following section cover searching for objects, downloading search results, and browsing the Inventory.

Searching overview

Use IBM® Storage Defender 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 Dell PowerMax Storage Inventory may include System Info, Volumes, Volume-Groups, Hosts, Host-Groups, Port-Groups, SRDF-Groups, and SRDF-Directors.

The Dell PowerFlex Storage Inventory may include System Info, Volumes, Hosts, Remote Consistency Groups, and Storage Data Replicators.

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 Defender Copy Data Management lets you quickly and easily locate every version of a file across your entire Enterprise. IBM® Storage Defender 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 information

[Searching for objects](#)

[Viewing object details](#)

[Viewing NetApp ONTAP file details](#)

[Finding and restoring a file](#)

[Downloading search results](#)

[Browsing inventory](#)

Searching for objects

Use the search feature to find objects on providers that are cataloged in IBM® Storage Defender 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 Defender 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 Defender 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 storage provider” on page 40](#) 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 Defender 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 262](#).
- 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 260](#).

Related information

[Downloading search results](#)

[Browsing inventory](#)

[Viewing object details](#)

[Viewing NetApp ONTAP file details](#)

[Search and filter guidelines](#)

[Select, sort, and reorder columns](#)

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 information

[Viewing NetApp ONTAP file details](#)

[Searching overview](#)

[Searching for objects](#)

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 125](#), and [“Creating a NetApp ONTAP File Inventory job definition” on page 127](#).
- Search for a file by selecting **NetApp > File** from the Search For filter on the Advanced Search dialog. See [“Searching for objects” on page 256](#).

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 262](#).

Related information

[Viewing object details](#)

[Searching overview](#)

[Searching for objects](#)

[Finding and restoring a file](#)

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 Defender 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 127](#).
- To enable NFS links from the **File Links** tab on an object's property pane, access IBM® Storage Defender 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 Defender 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 information

[Viewing NetApp ONTAP file details](#)

[Searching overview](#)

[Searching for objects](#)

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 Defender Copy Data Management. See [“Searching for objects” on page 256](#).

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 information

[Searching for objects](#)

[Viewing object details](#)

[Viewing NetApp ONTAP file details](#)

Browsing inventory

Use IBM® Storage Defender 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 Dell PowerMax Storage Inventory may include System Info, Volumes, Volume-Groups, Hosts, Host-Groups, Port-Groups, SRDF-Groups, and SRDF-Directors.

The Dell PowerFlex Storage Inventory may include System Info, Volumes, Hosts, Remote Consistency Groups, and Storage Data Replicators.

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 information

[Searching overview](#)

[Searching for objects](#)

[Viewing object details](#)

[Viewing NetApp ONTAP file details](#)

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**.

Report

The topics in the following section cover running and customizing reports, as well as individual report details.

Report Overview

IBM® Storage Defender 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 information

[Running a report](#)

[Creating a customized report](#)

[Editing a customized report](#)

[Downloading a Report](#)

[Creating a report job definition](#)

[Application reports](#)

[System management reports](#)

[File analytics reports](#)

[Protection compliance reports](#)

[Storage protection reports](#)

[Storage utilization reports](#)

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 storage provider” on page 40](#).
- Run a job immediately; sometimes referred to as on demand. See [“Start, pause, and hold a job session” on page 114](#).
- 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 111](#).

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 253](#).
- Save a report with customized parameters. See [“Creating a customized report” on page 266](#)

Related information

[Creating a customized report](#)
[Creating a report job definition](#)
[Downloading a Report](#)
[Report Overview](#)

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 storage provider” on page 40](#).

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 265](#).

Related information

[Editing a customized report](#)
[Running a report](#)

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 storage provider” on page 40](#).
- Run a job immediately; sometimes referred to as on demand. See [“Start, pause, and hold a job session” on page 114](#).
- Create a customized report. See [“Creating a customized report” on page 266](#).

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 265](#).

Related information

[Creating a customized report](#)
[Running a report](#)
[Downloading a Report](#)
[Report Overview](#)

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 265](#).

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 Defender 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 information

[Running a report](#)

[Creating a customized report](#)

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 265](#).

Related information

[Running a report](#)

[Downloading a Report](#)

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 information

[Application configuration report](#)

[Application RPO Compliance report](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

[Downloading a Report](#)

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 119](#).
- For Microsoft™-based application servers, ensure the vCenter containing the virtual machines are registered in IBM® Storage Defender Copy Data Management. Ensure that the virtual machine credentials for the virtual machines hosting the applications are configured. See [“Adding a credential” on page 87](#).

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 information

[Application reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 119](#).
- 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.

*The **Source** and **Destination** columns are deprecated in the IBM® Storage Defender Copy Data Management 2.2.26 release.

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.

*The **Source** and **Destination** columns are deprecated in the IBM® Storage Defender Copy Data Management 2.2.26 release.

Related information

[Application reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

System management reports

System Management Reports offer an in-depth view of the status of your IBM® Storage Defender 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 Defender 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 information

[Catalog summary report](#)

[Application configuration report](#)

[Job report](#)

[Job sessions report](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

[Downloading a Report](#)

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, Dell PowerMax, Dell PowerFlex, 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, Pure Storage FlashArray Objects, Dell PowerMax, and Dell PowerFlex 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 information

[System management reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

Configuration report

Review the IBM® Storage Defender Copy Data Management providers added to IBM® Storage Defender 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, Dell PowerMax Nodes, Dell PowerFlex 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.

Dell PowerMax Nodes

The following fields and corresponding data display in the Dell PowerMax Nodes section of the Configuration report:

Node (Site)

The name of the Dell PowerMax node and the associated site.

Model

The model number of the Dell PowerMax node.

Version

The Dell PowerMax version number.

Dell PowerFlex Nodes

The following fields and corresponding data display in the Dell PowerFlex Nodes section of the Configuration report:

Node (Site)

The name of the Dell PowerFlex node and the associated site.

Model

The model number of the Dell PowerFlex node.

Version

The Dell PowerFlex version number.

Related information

[System management reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

Job report

Review the available jobs in your IBM® Storage Defender 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 information

[System management reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

Job sessions report

Review sessions for jobs and their status in your IBM® Storage Defender 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 information

[System management reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 Defender 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 information

[File usage by owner report](#)

[Files by age report](#)

[Files by category report](#)

[Files by size report](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

[Downloading a Report](#)

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 Defender Copy Data Management Filewalker tool. See [“Creating a NetApp ONTAP File Inventory job definition” on page 127](#).

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 information

[File analytics reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 Defender Copy Data Management Filewalker tool. See [“Creating a NetApp ONTAP File Inventory job definition” on page 127](#).

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 information

[File analytics reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 information

[File analytics reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 Defender Copy Data Management Filewalker tool. See [“Creating a NetApp ONTAP File Inventory job definition” on page 127](#).

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 Defender 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 information

[File analytics reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 Defender 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 information

[File System RPO compliance report](#)
[IBM Storage Virtualize RPO compliance report](#)
[NetApp ONTAP protection usage report](#)
[NetApp ONTAP RPO compliance report](#)
[Pure Storage FlashArray RPO compliance report](#)
[Recovery points report](#)
[Unprotected Virtual Machines report](#)
[VMware RPO compliance report](#)
[Report Overview](#)
[Running a report](#)
[Creating a customized report](#)
[Downloading a Report](#)

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 121](#).
- Create and run a File System Backup job. See [“Creating a File System Backup job” on page 153](#).

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.

*The **Source** and **Destination** columns are deprecated in the IBM® Storage Defender Copy Data Management 2.2.26 release.

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 information

[Protection compliance reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 122](#).
- Create and run an IBM® Storage Virtualize Backup job. See [“Creating an IBM Storage Virtualize Backup job definition” on page 156](#).

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 information

[Protection compliance reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

NetApp ONTAP protection usage report

The NetApp ONTAP Protection Space Usage report displays the storage usage of IBM® Storage Defender 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 125](#).
- Create and run a NetApp ONTAP Backup job. See [“Creating a NetApp ONTAP Backup job definition” on page 169](#).

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 Defender 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 Defender 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 Defender 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 information

[Protection compliance reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 125](#).
- Create and run a NetApp ONTAP Backup job. See [“Creating a NetApp ONTAP Backup job definition” on page 169](#).

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 information

[Protection compliance reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 130](#).
- Create and run a Pure Storage Backup job. See [“Creating a Pure Storage FlashArray Backup job definition” on page 171](#).

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 information

[Protection compliance reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 information

[Protection compliance reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 131](#).
- All related storage providers must be added to IBM® Storage Defender 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 125](#).
- Create and run a VMware Backup job. See [“Creating a VMware Backup job definition” on page 174](#).

Use the Unprotected Virtual Machines report to answer questions such as:

- How many of my virtual machines are eligible for primary Storage Snapshot protection?
- How much storage space is used within a given time period?

Parameters

Use the following parameters to customize your report:

- vCenter
Multiple selections are supported.
- Power® State
Set the power state of virtual machines returned by the report. Power® states include Powered On, Powered Off, Suspended, or All. Multiple selections are supported. By default, this parameter is set to All.

The default report parameters report on unprotected virtual machines on all vCenters, in any power state.

Unprotected Virtual Machines

The following fields and corresponding data display in the Unprotected Virtual Machines section of the Unprotected Virtual Machines report for eligible virtual machines:

VM Name

The name of the virtual machine, preceded by an icon that displays the virtual machine's power state.

Location

The location of the virtual machine.

Hostname

The host node where the virtual machine resides.

Operating System

The operating system associated with the virtual machine

Provisioned Space

The amount of space on the datastore allocated for virtual disk files.

Datastores

The name of the associated datastore.

Related information

[Protection compliance reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 131](#).
- All related NetApp ONTAP storage providers must be added to IBM® Storage Defender 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 125](#).
- Create and run a VMware Backup job. See [“Creating a VMware Backup job definition” on page 174](#).

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 information

[Protection compliance reports](#)

[Report Overview](#)

[Running a report](#)

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 Defender 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 information

[NetApp ONTAP OSSV relationship status report](#)

[NetApp ONTAP overprotected volumes report](#)

[NetApp ONTAP Qtree protection status report](#)

[NetApp ONTAP underprotected volumes report](#)

[NetApp ONTAP volume protection status report](#)

[NetApp ONTAP transition dependency report](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

[Downloading a Report](#)

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 information

[Storage protection reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 information

[Storage protection reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 information

[Storage protection reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

NetApp ONTAP underprotected volumes report

The NetApp ONTAP Underprotected Volumes report displays volumes that are underprotected using NetApp ONTAP Snapshot, SnapVault, or SnapMirror software. Through your parameters selections, you can decide what constitutes an underprotected volume. Select the number of acceptable snapshots to reside on a volume, the lag time, and if a volume should also be SnapMirrored and SnapVaulted. If a volume has less protection than your parameter selections, it will be returned on the NetApp ONTAP Underprotected Volumes report.

Note: This report may discover storage systems that have not been cataloged in the database. These non-cataloged storage systems are discovered due to their replication relationships with other cataloged storage systems.

Use the NetApp ONTAP Underprotected Volumes report to answer questions such as:

- What are the names of my underprotected volumes?

- Which of my volumes are underprotected with less than 10 snapshots?

Parameters

Use the following parameters to customize your report:

- Storage Array
Multiple selections are supported.
- Acceptable Snapshots
- Lag time in days
- Is Volume SnapMirrored?
- Is Volume SnapVaulted?

The default report parameters display underprotected volumes with less than ten snapshots that are older than 7 days.

Underprotected Volumes

The following fields and corresponding data display in the Underprotected Volumes section of the NetApp ONTAP Underprotected Volumes report:

Volume

The name of the underprotected volume.

Storage Array

The node where the underprotected volume is located.

Location

The node where the volume resides.

Underprotection Reason

The reason the volume was returned by the report. For example, if the number of snapshots is lower than the defined No. of Acceptable Snapshots parameter, or if the snapshots have excessive lag times.

Related information

[Storage protection reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 information

[Storage protection reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 125](#).
- Create and run a NetApp ONTAP Backup job. See [“Creating a NetApp ONTAP Backup job definition” on page 169](#).

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 information

[Storage protection reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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, Dell PowerMax, Dell PowerFlex, 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 information

[IBM Storage Virtualize Consistency Groups Report](#)

[IBM Storage Virtualize Pools Report](#)

[IBM Storage Virtualize Volumes Report](#)

[Instant Disk Restore Volumes Report](#)

[NetApp ONTAP Aggregates Report](#)

[NetApp ONTAP LUNs Report](#)

[NetApp ONTAP Orphaned LUNs Report](#)

[NetApp ONTAP Quotas Report](#)

[NetApp ONTAP Snapshots Report](#)

[NetApp ONTAP Volumes Report](#)

[Pure Storage FlashArray Volumes Report](#)

[Storage Capacity Report](#)

[VM and storage mapping report](#)

[VMware Datastores Report](#)

[VMware LUNs Report](#)

[VMware Orphaned Datastores Report](#)

[VMware Orphaned LUNs Report](#)

[VMware VM Snapshot Sprawl Report](#)

[VMware VM Sprawl Report](#)

[VMware VM Storage Report](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

[Downloading a Report](#)

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 122](#).

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 Defender Copy Data Management where the consistency group resides.

Mapping Name

The mapping name IBM® Storage Defender 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 Defender Copy Data Management where the consistency group resides.

Mapping Name

The mapping name IBM® Storage Defender 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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 122](#).

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 Defender 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 Defender 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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 122](#).

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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 125](#).

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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 125](#).

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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 125](#).

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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 125](#).

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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 125](#).

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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 125](#).

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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 130](#).

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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

Storage Capacity Report

Report the storage capacity of your IBM® Storage Virtualize pools, Pure Storage FlashArray volumes, Dell PowerMax volumes, Dell PowerFlex 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 21](#), and [“Jobs overview” on page 114](#).
- 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 Defender 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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

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 131](#).

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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 131](#).

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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 131](#).

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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 131](#).

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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 131](#).

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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 131](#).

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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 131](#).

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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

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 131](#).

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 information

[Storage utilization reports](#)

[Report Overview](#)

[Running a report](#)

[Creating a customized report](#)

Maintenance

The topics in the following section cover maintenance information including updating the IBM® Storage Defender Copy Data Management appliance, logging on to the virtual appliance, and collecting logs for troubleshooting.

Maintenance overview

In most cases, IBM® Storage Defender 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 Defender 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 Defender Copy Data Management, VMware, and Linux® command-line usage. Maintenance tasks are performed in vSphere Client, through the IBM® Storage Defender Copy Data Management command-line, or through a web-base management console.

Maintenance tasks include collecting logs, updating the application, and reviewing the configuration of the virtual appliance.

Related information

[Log on to the virtual appliance](#)

[Setting the time zone](#)

[Collecting logs for troubleshooting](#)

[Modifying job log options](#)

[Managing the administrative console](#)

[Modifying the network settings](#)

[Uploading an SSL Certificate](#)

[Restoring a snapshot from a FlexGroup volume to another FlexGroup volume](#)

[LDAP username syntax](#)

Log on to the virtual appliance

Log on to the IBM® Storage Defender 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 Defender 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 information

[Collecting logs for troubleshooting](#)

[Managing the administrative console](#)

Setting the time zone

Access the IBM® Storage Defender Copy Data Management Administrative Console to set the time zone of the IBM® Storage Defender 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 Defender Copy Data Management logs and schedules will reflect the selected time zone. The selected time zone will also display on the IBM® Storage Defender 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 Defender Copy Data Management reside in time zones that differ from the IBM® Storage Defender Copy Data Management appliance, the instances can be synced to the time zone of the IBM® Storage Defender Copy Data Management appliance instead of the time zone of the instance. In some cases, a user may wish to schedule IBM® Storage Defender Copy Data Management jobs based on the IBM® Storage Defender 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 Defender Copy Data Management instances with the IBM® Storage Defender Copy Data Management appliance:

6. To sync time zones of IBM® Storage Defender Copy Data Management instances with the IBM® Storage Defender Copy Data Management appliance, complete the following steps:
 - a. Follow the procedure above to set the time zone of the IBM® Storage Defender Copy Data Management appliance.
 - b. Log in to the IBM® Storage Defender 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 Defender Copy Data Management appliance through the following commands:

```
service virgo stop
service virgo start
```

Related information

[Log on to the virtual appliance](#)

[Managing the administrative console](#)

Setting a Daylight Saving Time (DST) configuration

The Copy Data Management application may experience scheduling inconsistencies during Daylight Saving Time (DST) transitions. To prevent these inconsistencies, set the DST configuration to skip non-existent hours.

Procedure

By default, Copy Data Management does not skip hours that are lost during Daylight Saving Time (DST) transitions, which can cause scheduling inconsistencies. To prevent these issues, configure the application to skip the non-existent hours during DST changes.

Note: Issue all the `curl` commands from the command-line interface of the Copy Data Management appliance.

1. Create a session and obtain the session ID.

To authenticate and start a session, issue the following `curl` command from the command-line interface of the Copy Data Management appliance:

```
curl -k -s -X POST "https://localhost:8443/api/endeavour/session" -u
"<username>:<password>"
```

The command provides an output as follows:

```
{
  "sessionid": "<session-id>",
  "user": {
    .....
  }
}
```

Note the value of the session ID from the output for further use.

2. Retrieve the current DST configuration.

Issue the following command to retrieve the current DST configuration:

```
curl -s -k -X GET "https://localhost:8443/api/endeavour/preference/
pref.job.schedule.trigger.skipDayIfHourDoesNotExist" \
-H "Content-Type: application/json" \
-H "X-Endeavour-Sessionid: <session-id>"
```

Note: Replace `<session-id>` with the actual session ID obtained from the login API.

The command provides a response as follows:

```
{
  "name": "pref.job.schedule.trigger.skipDayIfHourDoesNotExist",
  "category": "Job",
  "type": "boolean",
  "value": {
    "value": null
  },
  ...
}
```

If the `value` attribute is `null`, the DST configuration is not enabled.

3. Enable DST configuration.

To enable DST configuration, issue the following command to set the `value` attribute to `true`.

```
curl -s -k -X PUT "https://localhost:8443/api/endeavour/preference/
pref.job.schedule.trigger.skipDayIfHourDoesNotExist" \
-H "Content-Type: application/json" \
-H "X-Endeavour-Sessionid: <session-id>" \
--data-raw '{
  "value": true
}'
```

Note: Replace `<session-id>` with the actual session ID obtained from the login API.

The request enables the DST configuration by setting the `pref.job.schedule.trigger.skipDayIfHourDoesNotExist` preference to `true`.

4. Restart the Virgo service by using the following steps:
 - a. Log in to the **Admin Console**.
 - b. Navigate to **System Actions > Restart IBM Storage Defender CDM**.
 - c. Select **Restart**.
 - d. Verify that the service is restarted.

What to do next

To verify that the DST configuration is enabled:

- Repeat [Step 2](#) to verify that the `value` attribute is set to `true`.

If you still see issues related to the processing of invalid hours:

- Verify the JSON format in the `curl` commands.
- Review the Copy Data Management logs for errors.

Collecting logs for troubleshooting

For troubleshooting the IBM® Storage Defender Copy Data Management application, IBM® Storage Defender 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 Defender 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 Defender Copy Data Management. Information included in the logs include the access time in EPOCH format, IBM® Storage Defender 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 Defender 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 Defender 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 information

[Log on to the virtual appliance](#)

[Monitor a job session](#)

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 information

[Log on to the virtual appliance](#)

[Monitor a job session](#)

Collecting the IBM® Storage Defender Copy Data Management logs from the virtual appliance

This procedure assumes IBM® Storage Defender 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 Defender Copy Data Management is deployed. See [“User administration and security management” on page 24](#).
- 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 Defender Copy Data Management logs from the virtual appliance, complete the following steps:

1. Log on to the IBM® Storage Defender 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 information

[Log on to the virtual appliance](#)

[Monitor a job session](#)

Troubleshooting policy-based high availability and policy-based replication errors

This topic describes common errors and their resolution.

Error in restoring to a host cluster when all hosts are not part of the same partition or are outside of it

Error

When you try to restore to a host cluster where all the hosts are not in the same partition or outside of it, you see the following error message in the job log:

```
Map LUNs: Mapping volume (Ora-fra-0) to host (xyz.host.abc.com) using protocols (FC) for ESX server (xyz.host.abc.com) Command mkvdiskhostmap -force -host "18" "14" failed with exception com.ibm.cdm.ibmsvc2provider.exception.IBMSVC2ExecCommandException: CMMVC1034E: The command failed because the host and the volume group the volume is part of the same storage partition or neither associated with a storage partition.
```

Resolution

The message is shown when you perform the restore operation and the destination host is part of a cluster and all the hosts are not added to the same partition. Therefore, when all the hosts in a cluster are in the same partition or outside of it, the restore operation works.

Error in backup job when the security scan server host is in a PBHA partition

Error

When the security scan server is on a host in a PBHA partition, you see the following error message in the job log:

```
Command mkvolumegroup -name "cdm_sscan_1031_1031_1754565653073_SLA_PBHA_SGC_SecurityScan_103" -type "thinclone" -fromsnapshotid "80" -partition "1" failed with exception com.ibm.cdm.ibmsvc2provider.exception.IBMSVC2ExecCommandException: CMMVC1243E: The command failed because the volume group is part of a storage partition that has a replication policy. Aug 7 16:58:03 2025 8 Failed to create snapshot clone of volume (fs7300a_dev1_rajup_iris_db) from snapshot (SLA_PBHA_SGC_SecurityScan_1031_1007_129_fs7300a_dev1_rajup_iris). Error = (StorageVolumeCloneException). Aug 7 16:58:03 2025 8 Failed to clone volume (fs7300a_dev1_rajup_iris_db). Exception (...serviceprovider.common.storage.exception.StorageVolumeCloneException) reports: (The command failed because the volume group is part of a storage partition that has a replication policy.)
```

Resolution

The message is displayed when the security scan server is deployed on an ESXi host which is in a PBHA partition. In IBM Storage FlashSystem 9.1.0, users cannot clone from a snapshot into a partition which has a PBHA policy. You need to move the security scan server on a host which is either outside of the partition or in a partition that does not have the PBHA policy enabled.

Error in restoring any application when the destination is part of a PBHA partition

Error

When you try to restore any application and select a destination within a PBHA partition, you see the following error message in job log:

```
Command mkvolumegroup -name "ECX_1011_1750739873197_1006_1007_173_-pbha-test" -type "thinclone" -fromsnapshotid "4" -partition "0" failed with exception com.ibm.cdm.ibmsvc2provider.exception.IBMSVC2ExecCommandException: CMMVC1243E: The command failed because the volume group is part of a storage partition that has a replication policy.
```

Resolution

The message is displayed when the destination of the restore operation is in a PBHA partition. To resolve the error, select a destination for restore that is not part of a PBHA partition.

Modifying job log options

The settings for job logs can be edited for the IBM® Storage Defender 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 Defender Copy Data Management appliance as the root user.
2. Stop the virgo service on the IBM® Storage Defender 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 Defender Copy Data Management appliance.

```
$ service virgo start
```

Updating global settings

You are able to modify a predefined set of configurable options in IBM® Storage Defender 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 Defender 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 Defender 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 **Job log purge** and **Job log purge offload** be enabled and is required when **Job log purge offload** is enabled.

IBM® Storage Defender Copy Data Management Server IP Address

Select the IP address that is assigned to your IBM® Storage Defender Copy Data Management appliance. The primary IBM® Storage Defender 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 Defender 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 Defender Copy Data Management will interact with if you have an IBM® Storage Defender 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 Defender 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 Defender 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 information

[Managing the administrative console](#)

Managing the administrative console

Log on to the Administrative Console to review the configuration of the IBM® Storage Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender 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 Defender Copy Data Management virtual appliance by changing the content of the configuration file and restarting Virgo.

6. Log in to the IBM® Storage Defender 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
```

Result

The heap of the Virgo component in IBM® Storage Defender Copy Data Management virtual appliance is updated to take advantage of increased memory and the number of concurrent processes is updated to optimize performance.

Upgrade IBM® Storage Defender Copy Data Management

This topic describes how to upgrade IBM® Storage Defender Copy Data Management.

Important: Internet access is required to upgrade the Copy Data Management appliance to version 2.2.28. If internet access is unavailable in your environment, please contact support for the necessary steps.

Important: Before you upgrade, take a VM snapshot of the CDM appliance from vCenter so that you have a stable snapshot to roll back to the previous state if necessary.

Note: When upgrading IBM® Storage Defender 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 upgrade applied is successful, but the upgrade cleanup process is failed. This message appears consistently when performing an upgrade to IBM® Storage Defender Copy Data Management, regardless of whether the Clean Process is failed or successful.

Important: You must cancel or move to permanent all **Resource Active** sessions before you upgrade to IBM® Storage Defender Copy Data Management 2.2.26 efix196 or a higher version from a previous version of CDM. Jobs from previous versions (prior to 2.2.26 efix195 versions) of IBM® Storage Defender 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.

Upgrading IBM® Storage Defender Copy Data Management appliance from 2.2.18 or later versions to 2.2.25 or later versions

Use this procedure to upgrade IBM® Storage Defender Copy Data Management appliance from 2.2.18 or later versions to 2.2.25 or later versions.

Before you begin

You can upgrade IBM® Storage Defender Copy Data Management to the latest version. If you are using an older version, upgrade to 2.2.18.0 or 2.2.19.0, and then to the latest version. For more information, see [“Upgrading IBM Storage Defender Copy Data Management appliance from 2.2.16 or 2.2.17 version to 2.2.18 or later versions” on page 347](#).

- Restart the virtual machine appliance before the upgrade.
- Ensure that IBM® Storage Defender Copy Data Management is installed and running. For more information, see [“Installing IBM Storage Defender Copy Data Management as a virtual appliance” on page 32](#).
- After IBM® Storage Defender Copy Data Management upgrades, you cannot rollback to a previous version without a virtual machine snapshot. Create a virtual machine snapshot of your environment before upgrading. If necessary, you can perform the virtual machine snapshot rollback to return to a previous version of IBM® Storage Defender Copy Data Management.

Important: Before upgrading IBM® Storage Defender Copy Data Management, you must change the root and administrator account passwords through SSH in the IBM® Storage Defender Copy Data Management appliance. SSH into the appliance with the username root or administrator and the default password. After authenticating, you will be prompted to change the password from the default.

Procedure

To upgrade your IBM® Storage Defender Copy Data Management appliance, complete the following steps:

1. Download the upgrade file from the 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**.
5. Click **Choose File**, browse for the upgrade file to upload to the appliance, then click **Upload Update Image**. The upgrade process begins after the upgrade image has been uploaded to the appliance.
6. After the upgrade 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 Defender Copy Data Management may be stored in your browser's cache. Clear your browser's cache before logging in to an upgraded version of IBM® Storage Defender Copy Data Management to ensure you are viewing the latest changes.

Remember:

- Upgrading the IBM® Storage Defender Copy Data Management appliance will require that the virtual machine be restarted for plugins to function properly.
- Ignore the **Clean process Failure** error while upgrading the IBM® Storage Defender Copy Data Management appliance.

What to do next

- If the upgrade process fails, encounters an error, or there is any interruption in the upgrade process prior to it completing, review the upgrade log files at the following location on the virtual machine:
/data/log/adminconsole
 - Collect the following log files:
 - update-manager.log
 - SCDMAdminConsole-upgrade-*.log
 - adminconsole.log.*.gz
 - OS-update-*.log
 - SCDMProduct-upgrade-*.log
- Additionally, collect the MongoDB upgrade logs (mongoUpgrade_*.log) from the following location:
/opt/scdm/tools/mongo-upgrade/logs/
- To reapply the upgrade, revert the virtual machine snapshot, which was created before the upgrade procedure, and then attempt another upgrade.

Related information

[Installing IBM Storage Defender Copy Data Management as a virtual appliance](#)

[Managing the administrative console](#)

Upgrading IBM® Storage Defender Copy Data Management appliance from 2.2.16 or 2.2.17 version to 2.2.18 or later versions

Use this procedure to upgrade IBM® Storage Defender Copy Data Management appliance version 2.2.16 or 2.2.17 to 2.2.18 or later versions.

Before you begin

To upgrade the IBM® Storage Defender Copy Data Management appliance to version 2.2.18 or later, make sure that the current IBM® Storage Defender Copy Data Management appliance version is either 2.2.16 or 2.2.17. If not, upgrade the IBM® Storage Defender Copy Data Management appliance previous version to 2.2.16, refer to [Upgrading IBM® Storage Defender Copy Data Management appliance from 2.2.15 or earlier versions to 2.2.16 or later versions](#).

Important: Before upgrading IBM® Storage Defender Copy Data Management, you must change the root account password through SSH in the IBM® Storage Defender 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 upgrade IBM® Storage Defender 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 Defender Copy Data Management virtual appliance.
2. Login to the IBM® Storage Defender 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 the upgrade script by issuing the following commands:

```
mkdir -p /mnt/SCDMiso
mount -o loop /tmp/<scdm ISO> /mnt/SCDMiso
cd /mnt/SCDMiso
./upgradeSCDM.sh
```

Note: After upgrading to 2.2.18 or later using the script, you must reboot the appliance.

Upgrading IBM® Storage Defender Copy Data Management appliance from 2.2.15 or earlier versions to 2.2.16 or later versions

Use this procedure to upgrade IBM® Storage Defender Copy Data Management appliance from 2.2.15 or earlier versions to 2.2.16 or later versions.

Before you begin

You can upgrade IBM® Storage Defender Copy Data Management to the most recent version. When you upgrade IBM® Storage Defender Copy Data Management, you must upgrade from no more than two previous versions (n-2) to the current version (n). If you are using an older version, you must upgrade at least to (n-2) version and then upgrade to the current version. When you upgrade to IBM® Storage Defender Copy Data Management 2.2.16.0, you must complete a new installation of IBM® Storage Defender Copy Data Management and then migrate from the previous version. This requires the catalog to be transferred from the previous version of IBM® Storage Defender Copy Data Management to the new deployment of IBM® Storage Defender 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 Defender Copy Data Management](#).

- Ensure IBM® Storage Defender Copy Data Management is installed and running. See [“Installing IBM Storage Defender Copy Data Management as a virtual appliance” on page 32](#).
- After IBM® Storage Defender Copy Data Management upgrades, it cannot rollback to a previous version without a virtual machine snapshot. Create a virtual machine snapshot of your environment before upgrading, and, if necessary, you can perform the virtual machine snapshot rollback to return to a previous version of IBM® Storage Defender Copy Data Management.

Important: Before upgrading IBM® Storage Defender Copy Data Management, you must change the root account password through SSH in the IBM® Storage Defender 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 upgrade process fails, encounters an error, or there is any interruption in the upgrade process prior to it completing, review the upgrade log files at the following location on the virtual machine:
/data/log/adminconsole
 - Collect the following log files:
 - update-manager.log
 - SCDMAdminConsole-upgrade-*.log
 - adminconsole.log.*.gz
 - OS-update-*.log
 - SCDMProduct-upgrade-*.log
- Additionally, collect the MongoDB upgrade logs (mongoUpgrade_*.log) from the following location:
/opt/scdm/tools/mongo-upgrade/logs/
- To reapply the upgrade, revert the virtual machine snapshot, which was created before the upgrade procedure, and then attempt another upgrade.

Related information

[Installing IBM Storage Defender Copy Data Management as a virtual appliance](#)

[Managing the administrative console](#)

Transferring the catalog from a previous version of IBM® Storage Defender Copy Data Management

The catalog is critical to the operation of IBM® Storage Defender Copy Data Management. When you upgrade IBM® Storage Defender Copy Data Management, you must transfer the catalog from the previous version of IBM® Storage Defender Copy Data Management to the current version of IBM® Storage Defender Copy Data Management.

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, in order to upgrade to IBM® Storage Defender Copy Data Management from the versions prior to 2.2.21, you must deploy a new instance of IBM® Storage Defender Copy Data Management and transfer the catalog from the previous version prior to 2.2.21 and restore it in the newly deployed version.

If you are updating the IBM® Storage Defender 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 Defender 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 115](#).
3. Log in to the Administrative Console of the existing IBM® Storage Defender 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. SSH to the existing IBM® Storage Defender Copy Data Management virtual appliance. Copy the contents of /data/backup to another location separate from the existing IBM® Storage Defender Copy Data Management virtual appliance that can be transferred when needed in a later step.
6. Deploy IBM® Storage Defender Copy Data Management. For information on deploying IBM® Storage Defender Copy Data Management, refer to [“Installing IBM Storage Defender Copy Data Management as a virtual appliance” on page 32](#).
7. Change the default passwords on the newly deployed IBM® Storage Defender 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 28](#).

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 Defender 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 Defender 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 Defender 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 115](#).
13. Re-import the source IBM® Storage Defender Copy Data Management license into the target IBM® Storage Defender Copy Data Management unless the source IBM® Storage Defender Copy Data Management and target IBM® Storage Defender Copy Data Management both have permanent licenses. If the source IBM® Storage Defender Copy Data Management license is permanent and the target IBM® Storage Defender Copy Data Management license is trial, the source IBM® Storage Defender Copy Data Management license should be imported into the target IBM® Storage Defender Copy Data Management.
14. Power® off the older IBM® Storage Defender 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.

Result

The catalog from the previous version of IBM® Storage Defender Copy Data Management is migrated to the newly deployed version of IBM® Storage Defender Copy Data Management.

Backup and Restore the Catalog

The Catalog Manager provides IBM® Storage Defender Copy Data Management administrators with the ability to backup and restore the IBM® Storage Defender Copy Data Management Catalog.

Before you begin

To restore the catalog, copy the backup .tar.gz catalog backup file on the Copy Data Management appliance:

1. Locate the backup .tar.gz catalog backup file on Copy Data Management appliance:
 - For 2.2.28 release, you can find the file in the /tmp/scdm/catalog/backup/ directory

- For previous releases, you can find the file in the `/data/backup/` directory
2. On the Copy Data Management appliance where you want to run the restore catalog, copy the file to the `/tmp/scdm/catalog/restore/` directory.

Procedure

To manage your IBM® Storage Defender 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 **IBM Storage Defender CDM** from the drop-down menu. Enter admin as a User and your admin password.
3. Click **Menu**, then select **Catalog Manager**.
4. Select Backup Catalog or Restore Catalog.
Backup Catalog: Click **Backup** to start the catalog backup. Your backup is saved to the `/tmp/scdm/catalog/backup/` directory

Backup Catalog considerations: IBM® Storage Defender Copy Data Management will be stopped while the Catalog is being backed up. The IBM® Storage Defender Copy Data Management user interface will not be accessible, and all running jobs will be aborted.

Restore Catalog: Click **Restore** to start the catalog restore.

Restore Catalog considerations: IBM® Storage Defender Copy Data Management will be stopped while the Catalog is being restored. The IBM® Storage Defender Copy Data Management user interface will not be accessible, and all running jobs will be aborted. All IBM® Storage Defender Copy Data Management snapshots created after the Catalog backup was run will be lost.

Related information

[Managing the administrative console](#)

Modifying the network settings

Access the IBM® Storage Defender 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 Defender 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 Defender 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 28](#).
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 Defender Copy Data Management IP or hostname change, complete the following steps:
Changes made to the IP or hostname of the IBM® Storage Defender 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 Defender Copy Data Management appliance.

SSH to each VADP proxy and enter the following commands:

- a. Enter the new IBM® Storage Defender Copy Data Management appliance IP or hostname to the ECX_HOST variable in /opt/ECX/bin/escvadb and save:

```
vi /opt/ECX/bin/escvadb
ECX_HOST=<new ECX IP or Hostname>
```

- b. Enter the new IBM® Storage Defender Copy Data Management appliance IP or hostname to the ECX_HOST variable in /etc/rc.d/init.d/escvadb and save:

```
vi /etc/rc.d/init.d/escvadb
ECX_HOST=<new ECX IP or Hostname>
```

- c. Restart the escvadb service:

```
service escvadb restart
```

Any associated VADP proxy servers should be now able to communicate with the IBM® Storage Defender Copy Data Management appliance using the updated IP address or hostname.

Related information

[Log on to the virtual appliance](#)

[Managing the administrative console](#)

Uploading an SSL Certificate

To establish secure connections in IBM® Storage Defender Copy Data Management, you must upload an SSL certificate through the web-based management console of the virtual machine where IBM® Storage Defender 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 Defender Copy Data Management.
- If you are uploading an LDAP SSL certificate, review LDAP syntax. See [“LDAP username syntax” on page 359](#).
- To import the HTTP certificate, enter the key store password in the password field. The default key store password is ecx-beta.

- Make sure that the CA certificate has the FQDN (Fully Qualified Domain Name) set in the SAN (Subject Alternative Name).

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 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](#)

[Register a provider](#)

[LDAP username syntax](#)

[User administration and security management](#)

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
```

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 Defender 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 Defender 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 information

[About the help system](#)

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 Defender 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 Defender 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 Defender Copy Data Management to use LDAP, and configuring IBM® Storage Defender Copy Data Management to use HTTPS.

Related information

[Documentation roadmap](#)

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 information

[Search for objects](#)

[Viewing object details](#)

[Viewing NetApp ONTAP file details](#)

[Creating a NetApp ONTAP File Inventory job definition](#)

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 information

[Viewing a provider](#)

[Editing a schedule](#)

[Editing a job definition](#)

[Monitor a job session](#)

[Searching overview](#)

[Running a report](#)

LDAP username syntax

LDAP form is used when setting up LDAP authentication in IBM® Storage Defender 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 information

[Uploading an SSL Certificate](#)

[User administration and security management](#)

Return code reference

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

Return Code (128+N)	Return Code Value
9	SIGKILL
10	SIGUSR1
11	SIGSEGV
12	SIGUSR2
13	SIGPIPE
14	SIGALRM
15	SIGTERM
16	SIGSTKFLT
17	SIGCHLD
18	SIGCONT
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

Return Code (128+N)	Return Code Value
51	SIGRTMAX-13
52	SIGRTMAX-12
53	SIGRTMAX-11
54	SIGRTMAX-10
55	SIGRTMAX-9
56	SIGRTMAX-8
57	SIGRTMAX-7
58	SIGRTMAX-6
59	SIGRTMAX-5
60	SIGRTMAX-4
61	SIGRTMAX-3
62	SIGRTMAX-2
63	SIGRTMAX-1
64	SIGRTMAX

Related information

[Creating a script job definition](#)

Tuning external parameterized configurations

The topic describes properties for tuning the external parameterized configurations for various scanning operations.

Sentinel NFS tuning configuration

The following table describes the Sentinel NFS tuning related configuration properties that are available in the /opt/ECX/virgo/repository/ecx-usr/com.syncsort.dp.xsb.serviceprovider.properties file.

Table 19: Sentinel NFS tuning related configuration properties

Property name	Default value	Applicable workflow	Description
protection.application.scan.nfs.rsize	0	AIX based application scanning	Configures thesizeNFS mount command option, which is executed to mount the NFS share on a security scan server host during the AIX scanning workflow.
protection.application.scan.nfs.wsize	0	AIX based application scanning	Configures thewsizenfs mount command option, which is executed to mount the NFS share on the Cybersense host during the AIX scanning workflow.
protection.application.scan.nfs.nconnect	0	AIX based application scanning	Configures thenconnectNFS mount command option, which is executed to mount the NFS share on the Cybersense host during the AIX scanning workflow.
protection.application.scan.nfs.actimeo	0	AIX based application scanning	Configures theactimeoNFS mount command option, which is executed to mount the NFS share on the Cybersense host during the AIX scanning workflow.
protection.application.scan.nfs.proxy.highSpeedNetworkIPs	0	AIX based application scanning	Accepts one or more comma-separated alternative proxy host IP addresses (high-speed network IP) for the AIX scanning workflow.
protection.application.scan.nfs.indexEngineHostIP	0	AIX based application scanning	Accepts an alternative index engine host IP address (high-speed network IP) for the AIX scanning workflow.

Sentinel scanning configuration

The following table describes the Sentinel scanning related general properties that are available in the `/opt/ECX/virgo/repository/ecx-usr/com.syncsort.dp.xsb.serviceprovider.properties` file.

Table 20: Sentinel scanning related general properties

Property name	Default value	Relevant workflow	Description
<code>protection.application.scan.scanISCSISession</code>	false	VMware VM scanning	Accepts a Boolean value, when set to true, the backup job does the SCSI scanning multiple times to counter the network slowness in the VMware scanning.
<code>protection.application.AIX.proxy.unmount.timeout.seconds</code>	600	AIX based application scanning	Configures a timeout in seconds for the unmount command that is executed on the AIX host during the scanning workflow execution.
<code>protection.application.scan.fullNFSRestriction</code>	true	AIX based application Scanning	Tunes the NFS exports restriction in the AIX based scanning workflow. If set to true, the option sets the full restriction and the NFS shares that are created and exported to the index engine host only. If set to false, the option creates a global NFS share to all the IP addresses available in the network.

Remote script execution configuration

To execute scripts on a Linux server custom remote directory, you can configure Copy Data Management with the remote directory location where the script is expected to be uploaded and then executed. You can specify the remote directory for executing custom scripts by using the `com.catalogic.ecx.remotexecutor.properties` configuration file. The file is available in the `/opt/virgo/repository/ecx-usr` directory.

Table 21: Remote script execution configuration property

Property name	Default value	Relevant workflow	Description
<code>remote.config.script.destinationLocation</code>	/tmp	Jobs with custom scripts	Specifies the directory location where the custom scripts are uploaded and then executed as remote scripts.

Clean rate configuration

Maintenance jobs might take longer to complete because FlashCopies often take more than 10 minutes to transition into the Stopped state. Copy Data Management monitors the FlashCopy status for up to 10 minutes. If the FlashCopy state is not changed to Stopped within 10 minutes, Copy Data Management proceeds to the next FlashCopy without deleting the current one.

To improve the clean-up rate before stopping FlashCopy on IBM Storage Virtualize, you can configure the clean rate up to 100. You can enable and configure the FlashCopy clean rate before halting and deleting the FlashCopy during the maintenance window of 30 minutes.

Table 22: Clean rate configuration properties

Property name	Default value	Relevant workflow	Description
<code>ibmsvc.flashcopy.stop.cleanrate.update</code>	false	Maintenance workflow	Accepts a Boolean value to enable or disable the clean rate configuration.
<code>ibmsvc.flashcopy.stop.cleanrate.update.value</code>	100	Maintenance workflow	Accepts a numeric integer value. The property is applicable only when the <code>ibmsvc.flashcopy.stop.cleanrate.update</code> property is set to true. You can provide a value in the range of 1 to 100.

In-backup mode configuration

You can enable the in-backup mode condense operations for IBM Storage Virtualize, IBM Storage Virtualize for Snapshot, and Dell PowerMax, backup jobs for maintenance.

You can enable the in-backup mode condense operation for the respective storage platforms by updating the following properties.

Table 23: In-backup mode configuration properties

Property name	Default value	Relevant workflow	Description
pref.job.condense.inBackup.ibm_svc	false	Backup jobs based on IBM Storage Virtualize	Accepts a Boolean value as true or false. If set to true, enables the in-backup condense mode for backup jobs based on IBM Storage Virtualize.
pref.job.condense.inBackup.vmware.ibm_svc	false	VM backup jobs based on IBM Storage Virtualize	Accepts a Boolean value as true or false. If set to true, enables the in-backup condense mode for VM backup jobs based on IBM Storage Virtualize.
pref.job.condense.inBackup.application.ibm_svc	false	App backup jobs based on IBM Storage Virtualize	Accepts a Boolean value as true or false. If set to true, enables the in-backup condense mode for app backup jobs based on IBM Storage Virtualize.
pref.job.condense.inBackup.ibm_svc2	false	Backup jobs based on IBM Storage Virtualize for Snapshot	Accepts a Boolean value as true or false. If set to true, enables the in-backup condense mode for backup jobs based on IBM Storage Virtualize for Snapshot.
pref.job.condense.inBackup.vmware.ibm_svc2	false	VM backup jobs based on IBM Storage Virtualize for Snapshot	Accepts a Boolean value as true or false. If set to true, enables the in-backup condense mode for VM backup jobs based on IBM Storage Virtualize for Snapshot.
pref.job.condense.inBackup.application.ibm_svc2	false	App backup jobs based on IBM Storage Virtualize for Snapshot	Accepts a Boolean value as true or false. If set to true, enables the in-backup condense mode for app backup jobs based on IBM Storage Virtualize for Snapshot.
pref.job.condense.inBackup.dellpowermax	false	Backup jobs based on Dell PowerMax	Accepts a Boolean value as true or false. If set to true, enables the in-backup condense mode for backup jobs based on Dell PowerMax.
pref.job.condense.inBackup.vmware.dellpowermax	false	VM backup jobs based on Dell PowerMax	Accepts a Boolean value as true or false. If set to true, enables the in-backup condense mode for VM backup jobs based on Dell PowerMax.
pref.job.condense.inBackup.application.dellpowermax	false	App backup jobs based on Dell PowerMax	Accepts a Boolean value as true or false. If set to true, enables the in-backup condense mode for app backup jobs based on Dell PowerMax.

- For versions till 2.2.24:
 - Go to the `/opt/ECX/virgo/repository/ecx-usr/preferences/` directory on the Copy Data Management appliance and change the value of properties as needed.
- For versions 2.2.25 and later:
 - Connect to Copy Data Management by using SSH and issue the command as shown in the following syntax:

```
curl https://<CDM-IP>:8443/api/endeavour/preference/<PROPERTY_NAME> -X PUT -H
'Accept: application/json' -H 'Content-Type: application/json' -H 'X-Endeavour-
Sessionid:<SESSION-ID>' --data-raw ' {"value":true}' --insecure
```

Repeat the command for each property you want to update.

Note: You must update the Copy Data Management instance IP address and the session ID in the command syntax according to your instance. To disable the in-backup mode, use the same command syntax and set the Boolean value to `false` instead of `true`.

Replication timeout configuration

During a backup for IBM Storage job by using the Global Mirror Change Volumes policy, Copy Data Management checks the `freezeTime` value 20 times at 30-second intervals and waits for approximately 10 minutes before you create a FlashCopy.

If you want Copy Data Management to wait longer than 10 minutes before creating the FlashCopy, you can increase the timeout by modifying the `protection.ibm_svc.replication.timeout` setting in the `com.syncsort.dp.xsb.serviceprovider.properties` file, which is located in the `/opt/virgo/repository/ecx-usr/` directory.

Table 24: Replication timeout configuration property

Property name	Default value	Relevant workflow	Description
<code>protection.ibm_svc.replication.timeout</code>	0	Backup jobs with GMCV sub policy for IBM Storage	Accepts a positive numerical integer as the number of polling cycles.

Oracle agent code execution timeout configuration

During Oracle backup and restore workflows, Oracle agent code execution timeout is configured for two hours by default. If the agent code execution exceeds the default two-hours limit, the operation times out.

Copy Data Management provides an externally configurable parameter, which defines the maximum duration allowed for Oracle agent code execution for Oracle workflows. You can set an appropriate timeout value to prevent agent code execution timeout.

The `sshTimeout` property is used to configure the Oracle agent code execution timeout. You can add the property in the `com.catalogic.ecx.deploy.vmware.ecxvmdeployer.json` file, which is located in `/opt/ECX/virgo/repository/ecx-usr` directory.

The following table describes the Oracle agent code execution timeout configuration property.

Table 25: Oracle agent code execution timeout configuration property

Property name	Default value	Relevant workflow	Description
<code>sshTimeout</code>	-	Oracle backup and restore workflows	Configure a timeout (in milliseconds) for Oracle agent code execution during the Oracle backup and restore workflows.

Example

Configuration example:

The default timeout is 7,200,000 milliseconds (2 hours). To increase it to 4 hours, add the `sshTimeout` property in the file and set the value to 14,400,000 milliseconds:

```
"dataLogPath":"/data/log/ecxdeployer",
"ecxGuestToolsPath":"/opt/ECX/guesttools",
"windowsGuestToolsWorkingDir":"c:
ProgramData",
-----
"transferTimeout":120000,
"rubyPath":"/usr/bin/ruby",
"sshTimeout":14400000
```

The guestapps agent configuration properties

You can configure the Copy Data Management guestapps agent properties in the `guestapps.conf` file, which is located in the `/etc` directory of the application host.

Note: If the `guestapps.conf` is not available in the `/etc` directory, create a file with the same name and add the property.

The following table describes the guestapps agent properties that you can configure for Copy Data Management.

Table 26: Configuration properties of guestapps agent			
Property name	Default value	Relevant workflow	Description
<code>sqlCommandTimeout</code>	300	Oracle backup and restore workflows	Configure a timeout (in seconds) for all Oracle SQL commands that are executed during Oracle backup and restore workflows by updating the <code>sqlCommandTimeout</code> property.

Important: Restart the Copy Data Management appliance to reflect the changes in properties. Use the following steps to restart the Copy Data Management appliance:

1. Use SSH to connect to the Copy Data Management appliance as the root user.
2. Stop the Copy Data Management appliance by using the following command:

```
systemctl stop virgo
```

3. Go to the required directory and open the file in the edit mode.
4. Update the property value and save the file.
5. Start the Copy Data Management by using the following command:

```
systemctl start virgo
```

Latest blog posts

Refer to blog posts to learn more about IBM® Storage Defender Copy Data Management and IBM® Storage Defender Sentinel.

- [Implementing PBR/PBHA backups with IBM Defender CDM](#)
- [Backup & Restore an SAP HANA database in an HSR Cluster](#)

- [Elevating Data Protection: New Capabilities in IBM Storage Defender Sentinel 1.1.11](#)
- [Detect ransomware threats on VMware VMs with Storage Copy Data Management](#)
- [Using IBM Storage Defender CDM for VMware - FC based Sentinel scanning](#)

Related information

[Creating a script job definition](#)

Frequently asked questions

The following are answers to frequently asked questions related to IBM® Storage Defender Copy Data Management functionality. The questions and answers are organized by deployment, resources, connectivity, cataloging, operation, and control topics.

Deployment

How is IBM® Storage Defender Copy Data Management distributed?

In most cases, IBM® Storage Defender Copy Data Management is distributed as a virtual appliance through an OVF template.

How do I configure IBM® Storage Defender Copy Data Management out-of-box?

See the topics [“Deployment checklist” on page 27](#), and [“Installing IBM Storage Defender Copy Data Management as a virtual appliance” on page 32](#).

What are the requirements of the datastores used for the hard disks? What types of VMware datastores are supported?

The type of datastore on which IBM® Storage Defender Copy Data Management is deployed is transparent to IBM® Storage Defender Copy Data Management.

Can the virtual appliance hard disks be thin provisioned?

Yes. This is a function of the virtual appliance, and can be set during IBM® Storage Defender Copy Data Management installation. Better performance can be achieved with thick provisioning of the appliance.

Why is there a delay from when the machine boots to login?

When you boot the machine, several processes occur including:

- Operating system and network connections initiate.
- Dependencies are scanned and resolved.

The more heavily loaded the ESX server is, the longer the boot process might take.

What are the default IBM® Storage Defender Copy Data Management user names and passwords?

When logging on to IBM® Storage Defender Copy Data Management for the first time, the default user name is admin and the default password is password. You will be prompted to reset the default password.

When logging on to the management console of the virtual machine, the default user name is administrator and the default password is ecxadLG235.

How is the root password secured?

You are prompted to change the root password on the first root login.

Resources

Can the disks be increased dynamically?

IBM® Storage Defender Copy Data Management data volumes can be expanded if necessary with the approval of Technical Support.

What resource can I add to improve IBM® Storage Defender Copy Data Management performance?

Increasing memory should help improve performance.

How much of the virtual machine default configuration can be modified?

Parameters such as network, CPU, and memory can be configured at the virtual machine level, but adjusting to below the default levels may impact performance. For default requirements, see the topics [System requirements](#), and [“Installing IBM Storage Defender Copy Data Management as a virtual appliance” on page 32](#).

Is it possible to install proprietary software, such as antivirus software, on the virtual appliance?

It is not recommended to install third party applications on the virtual appliance without approval from Technical Support.

Can I access the IBM® Storage Defender Copy Data Management user interface remotely?

Yes. The IBM® Storage Defender Copy Data Management user interface is browser based. Supported browsers and the URL are described in the topics [System requirements](#), and [“Starting IBM Storage Defender Copy Data Management” on page 35](#).

What ports are needed to access the IBM® Storage Defender Copy Data Management user interface?

To access IBM® Storage Defender Copy Data Management, appropriate ports need to be opened through the firewall. For details, see the topic [“Starting IBM Storage Defender Copy Data Management” on page 35](#), and [“User administration and security management” on page 24](#).

What operating system is IBM® Storage Defender Copy Data Management built on?

CentOS is the operating system on the IBM® Storage Defender Copy Data Management virtual appliance.

Is Java™ used as part of the IBM® Storage Defender Copy Data Management appliance?

Yes. However, OpenJDK is used as opposed to JRE.

Is JavaScript™ required for accessing IBM® Storage Defender Copy Data Management?

A browser that supports JavaScript™ is required. Supported browsers and the URL are described in the topics [System requirements](#).

Do the IBM® Storage Defender Copy Data Management cataloging and reporting functions impact the performance of the registered storage systems?

The cataloging function is built on technology that is designed to run as low priority on the storage system and automatically adjust itself to give top priority to primary workload operations. The reporting functions do not impact registered storage systems as they run on the IBM® Storage Defender Copy Data Management virtual appliance.

Connectivity

How does IBM® Storage Defender Copy Data Management connect to NetApp ONTAP storage systems?

IBM® Storage Defender Copy Data Management connects to NetApp ONTAP storage systems through HTTPS or HTTP.

Is the network traffic secure?

Network traffic between the IBM® Storage Defender Copy Data Management virtual appliance and the the IBM® Storage Defender Copy Data Management user interface is secured using HTTPS protocol. Network traffic between the IBM® Storage Defender Copy Data Management virtual appliance and an external resource, such as a NetApp storage system, a vCenter, or an LDAP server, uses either HTTPS or HTTP protocol, which is decided by the System Administrator when registering the resource.

Does IBM® Storage Defender Copy Data Management work with storage vendors other than IBM®, and NetApp?

IBM® Storage Defender Copy Data Management software works with IBM®, and NetApp storage and VMware infrastructure.

IBM® Storage Defender Copy Data Management provides Backup and Restore support for customers with VMware leveraging heterogeneous storage, extending Backup use cases to VMware on mixed storage.

For the search and reporting features of IBM® Storage Defender Copy Data Management, the VMware environment can use any storage; it does not have to be IBM®, or NetApp. Therefore, IBM® Storage Defender Copy Data Management provides visibility and insight into VM information across any storage device.

Does IBM® Storage Defender Copy Data Management work with volumes that have non-Windows™ file systems?

Yes. IBM® Storage Defender Copy Data Management catalogs NFS and CIFS files residing on NetApp ONTAP volume snapshots. Linux®/Unix files are stored using NFS and Windows™ files are stored using CIFS protocol.

Does IBM® Storage Defender Copy Data Management work with SnapManager data?

IBM® Storage Defender Copy Data Management catalogs the meta-data on LUNs created by SnapManager for SQL Server and Exchange. File level granularity of content hosted inside these LUNs is not available.

Is IBM® Storage Defender Copy Data Management software SNMP compliant?

Not at this time.

Cataloging

If you add a vCenter into an Inventory job definition, does that automatically discover all the ESX servers within that vCenter?

Yes. Once cataloged, view available VMware resources through the Inventory browse function on the Search tab.

How many storage systems can be cataloged?

IBM® Storage Defender Copy Data Management can catalog any number of storage systems and is only limited by the data disk in its delivered configuration.

When a catalog job runs, is it a full catalog job each time?

Yes. A full catalog job, not an incremental, is run each time.

Why is it that sometimes many jobs and tasks are marked with Waiting indicators on the Jobs tab?

The number of operations in progress on the Jobs tab varies depending on the number of jobs currently running. IBM® Storage Defender Copy Data Management controls the number of jobs allowed to run. When the number of jobs exceeds the value defined by IBM® Storage Defender Copy Data Management, jobs marked with Waiting indicators display on the Jobs tab. IBM® Storage Defender Copy Data Management also controls the number of job tasks to run simultaneously for a given job when multiple jobs are running.

When does the cataloged data get cleaned up?

After a certain number of runs for a given job, older objects for that job are purged from the Catalog. This retention parameter is set when the job is defined.

The Maintenance job removes resources and associated objects created by IBM® Storage Defender 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 Defender Copy Data Management catalog, and removes related snapshots.

See the topic [“Maintenance job” on page 253](#).

Operation

How do I protect and recover the IBM® Storage Defender Copy Data Management appliance itself?

Backing up the IBM® Storage Defender Copy Data Management appliance regularly is a critical operation. For more information, contact Technical Support.

Can I restore an individual file by using the Search window?

Yes. To restore a file on a Windows™ machine, search for the file in IBM® Storage Defender Copy Data Management and discover its location through the file's properties pane. See the topic [“Finding and restoring a file” on page 262](#).

The recommended Best Practice for protection/recovery is to use Backup and Restore jobs in IBM® Storage Defender Copy Data Management.

Why is it that when I select Hide Duplicates when performing an advanced search, some duplicate objects still display in the search results pane?

In some cases, the name of a returned object on the search results pane may be the same as another object, however the resources where the objects reside is different. Review the file properties of the objects by selecting their names on the search results pane to view the differences between the returned entries.

Can I generate customized reports?

IBM® Storage Defender Copy Data Management provides a set of predefined reports that can be customized through the use of parameter selection.

How are IBM® Storage Defender Copy Data Management logs collected for troubleshooting?

There are two approaches for downloading logs. Download logs from the Support menu or access the IBM® Storage Defender Copy Data Management appliance through a command prompt. The first approach is simpler and generally sufficient. The second approach produces a more comprehensive set of logs. See the topic [“Collecting logs for troubleshooting” on page 339](#).

Is audit tracking provided?

An audit log displaying IBM® Storage Defender Copy Data Management activity is available through the **Support** menu. Click the arrow next to the Support icon, then click **View Audit Log**.

Backup/Restore Jobs

For a Backup job, how do I update the retention after a job has run?

Open the existing job definition, click Snapshot in the workflow pane, and update the Keep Snapshots parameter. The retention policy changes to the supplied value when the job is next run.

To what extent do the IBM® Storage Defender Copy Data Management Backup and Restore functions impact the performance of NetApp storage systems?

The Backup and Restore functions employ technologies such as Snapshot and FlexClone that are designed to be low-impact on the NetApp storage systems. Generally, users should observe little unexpected performance impact on the storage systems.

For a Backup job, can I exclude swap partitions?

No. The lowest granularity of protection is a virtual machine.

Related information

[Oracle database support FAQ](#)

[Microsoft SQL Server Support FAQ](#)

Oracle database support FAQ

The following are answers to frequently asked questions related to IBM® Storage Defender Copy Data Management Oracle database support functionality.

What is Oracle CDM? How does it help solve my challenges?

Database administrators working extensively with Oracle are challenged when faced with mission-critical use cases such as Backup, Recovery, DevOps, and Business Analytics. This is especially true given that their Oracle databases have expanded in size and number over time, and that the databases need to be up and running 24x7x365.

Oracle DBAs struggle with the following:

- Backups are slow, complex and need constant management
- Backup process slows down the production servers
- Recoveries are slow and complex
- Repurposing App consistent backups (clones) for DevOps and Business Analytics is slow, complex and storage inefficient
- Lack of automation exists for producing quick and secure clones required to accelerate DevOps
- Copy sprawl problems occur due to no central catalog of copies
- Unable to meet organization stringent RPO and RTO requirements

IBM® Storage Defender Copy Data Management simplifies Oracle database copy management by enabling administrators to orchestrate application-consistent copy creation, cloning and recovery in minutes, instead of hours or days. IBM® Storage Defender Copy Data Management copy management leverages the advanced snapshot and replication features of the underlying storage platform to rapidly create, replicate, clone, and restore copies of Oracle databases in the most efficient way possible, in both time and space. IBM® Storage Defender Copy Data Management enables you to focus on the backup and recovery requirements of your business rather than the technical details of the underlying storage platforms.

IBM® Storage Defender Copy Data Management is an intelligent copy data management solution that delivers end-to-end automation, orchestration, and self-service functionality for your Oracle environment through a comprehensive and scalable Inventory. With the self-service features of IBM® Storage Defender Copy Data Management, your users are empowered to create clones on demand, freeing DBAs, while at the same time offering the advanced recovery features needed for Oracle environments.

IBM® Storage Defender Copy Data Management Oracle Copy Data Management solution supports the following Oracle deployment modes:

- Single instance – a single instance running on a single server accessing a database
- RAC (Real Application Clusters) leveraging ASM – more than one instance running on multiple servers are accessing a database simultaneously
- ASM (Automatic Storage Management) – Oracle's own volume manager and cluster filesystem that is optimized for Oracle Database features.

Deployment and Registration

Do I need to deploy any additional agents to protect Oracle standalone or Oracle RAC servers?

IBM® Storage Defender Copy Data Management for Oracle is delivered as a VMware OVA that is easily deployed on demand in a matter of minutes. Once deployed, you simply register your Oracle servers with appropriate credentials and then let IBM® Storage Defender Copy Data Management discover the rest. IBM® Storage Defender Copy Data Management eliminates the complexity of manually deploying and maintaining application agents on Oracle servers. A lightweight application-aware component is automatically injected to the required Oracle servers on demand and automatically updated to the latest version if required.

Oracle Backup creation workflow

App consistent Oracle database backup creation (local and remote) step by step

IBM® Storage Defender Copy Data Management auto-discovers databases and enables copies only of eligible databases. To be eligible for IBM® Storage Defender Copy Data Management backup, the Oracle database needs to be residing on a supported storage platform. With IBM® Storage Defender Copy Data Management, application owners do not need to be concerned about storage infrastructure.

IBM® Storage Defender Copy Data Management creates application-consistent Oracle database copies without the need to build and maintain complex RMAN scripts.

A typical IBM® Storage Defender Copy Data Management Oracle database backup creation workflow consists of the following steps:

- Auto-inject a lightweight component into the standalone Oracle Server(s) or one of the Oracle RAC server node(s) running a database instance to be copied
- Discover storage volume mapping to selected Oracle database(s) and logs
- Place the Oracle database in hot backup mode
- Automatically create a consistency group for related storage volumes
- Create an application-consistent backup of the consistency group
- Take the Oracle database out of hot backup mode (typically within a few seconds of entering the hot backup mode)
- Optionally create log copies into the specified mount points
- Optionally create selected masked copies using masking tools for secure DevOps use
- Catalog Oracle copies in Inventory and optionally record details in RMAN recovery catalog
- Optionally replace application-consistent backup to remote location leveraging storage replication
- Clean up auto-injected components from Oracle server node

IBM® Storage Defender Copy Data Management creates and uses in-place copies, so no data is physically moved. Replication to off-host storage is performed by leveraging storage replication, which reduces the amount of impact on Oracle Servers and Databases. IBM® Storage Defender Copy Data Management generated application-consistent copies are both space and time efficient. With the same ease, a DBA can automate the creation of remote copies for DR use cases.

Does IBM® Storage Defender Copy Data Management Oracle solution leverage the Storage Consistency Group feature?

The storage consistency group feature allows storage administrators to take a snapshot of database applications where the data is spread across multiple volumes to maintain consistency across all volumes.

In a typical Oracle Database, the data is spread across different volumes for better IO performance and availability. IBM® Storage Defender Copy Data Management Oracle application-consistent backup creation ensures that appropriate consistency groups are automatically created to maintain consistency across all related volumes.

What level of selection granularity is supported for Oracle backup workflows?

IBM® Storage Defender Copy Data Management Oracle backup workflows support copy selection at the following levels:

- One or more Oracle home locations
- One or more databases
- One or more container databases for Oracle 12c

Why can I not select some of the databases for protection in an Oracle backup workflow?

You cannot select a database if it is not eligible for protection. Hover your cursor over the database name to view the reasons the database is ineligible, such as that the database files, control files, or redo log files are stored on unsupported storage.

Will IBM® Storage Defender Copy Data Management auto discover newly added databases and automatically protect them?

If you select the parent Oracle Home in a Backup job definition, all databases under it are protected. If a new database is added under the home, it will be automatically protected once it is cataloged. Discovery and cataloging of new Oracle databases occurs as part of regularly scheduled Oracle Inventory job.

Does the Oracle database need to be on supported storage for IBM® Storage Defender Copy Data Management?

Yes, the IBM® Storage Defender Copy Data Management Oracle solution leverages storage snapshots for database protection. All databases, database files, control files and redo logs must be on supported storage systems for it to be eligible for protection.

Does IBM® Storage Defender Copy Data Management leverage the Oracle 12c Storage Snapshot feature?

This new feature of Oracle 12c enables you to take a storage snapshot of your database without needing the database to enter BACKUP mode. In Oracle, when you need to recover, you can use a point in time of the snapshot. You can roll forward by using the database archive logs, and use this snapshot feature to recover part or all of the database. IBM® Storage Defender Copy Data Management fully supports this feature starting in IBM® Storage Defender Copy Data Management 2.5.1.

Does IBM® Storage Defender Copy Data Management support protection of Offline Databases?

Databases in offline mode are not automatically included during backup workflows, unless they share storage volumes with a selected database that is active. IBM® Storage Defender Copy Data Management marks the offline databases and does not present them for Oracle Restore Workflows. Users may be able to retrieve these database files as flat files and perform application mount outside of IBM® Storage Defender Copy Data Management.

Does IBM® Storage Defender Copy Data Management support protection of Oracle databases not running in Archive Log mode?

Yes, protection of Oracle Databases running in NOARCHIVELOG mode are now supported for both Inventory and Backup use cases. You can recover database to the point of the most recent snapshot. PIT recoveries are not supported for databases running in NOARCHIVELOG mode.

Does IBM® Storage Defender Copy Data Management support protection of Oracle databases using pFile (text initialization parameter files)?

Yes, IBM® Storage Defender Copy Data Management now supports IBM® Storage Defender Copy Data Management backup of databases started through pFile in addition to spFile.

Oracle Archive log management

Does IBM® Storage Defender Copy Data Management support archive log backup and log management?

Oracle DBMS creates database transaction logs as part of its operation. Oracle databases can run in the following logging modes:

- NOARCHIVELOG mode – In NOARCHIVELOG mode, no transaction logs are created, and there is no capacity to run point-in-time recovery or online backups. This is the default.
- ARCHIVELOG mode – In ARCHIVELOG mode, the database makes copies of all online redo logs after they are filled. These copies are called archived redo logs. The archived redo logs are created via the ARCH process. The ARCH process copies the archived redo log files to one or more archive log destination directories. These saved archived logs are used for point-in-time recovery.

IBM® Storage Defender Copy Data Management provides you with an option for archive log files processing:

- Enable archive log backup (Recommended)
- Use existing archive log (Default)

IBM® Storage Defender Copy Data Management greatly simplifies archive log protection. If a user chooses to protect archive logs, IBM® Storage Defender Copy Data Management enables continuous backup of archive logs to a specified destination providing the lowest RPO (transaction level recoveries).

IBM® Storage Defender Copy Data Management automatically discovers the location where Oracle writes archived logs. If this location resides on storage from a supported vendor, IBM® Storage Defender Copy Data Management can protect it. If the existing location is not on supported storage, or if you wish to create an additional backup of database logs, enable the Create Additional Archive Log Destination option in the Oracle Backup job definition, then specify a path that resides on supported storage. When enabled, IBM® Storage Defender Copy Data Management configures the database to start writing archived logs to this new location in addition to any existing locations where the database is already writing logs. If multiple databases are selected for backup, then each of the servers hosting the database must have their destination directories set individually.

Can I specify a retention period for backed up archive logs within IBM® Storage Defender Copy Data Management?

If the Create Additional Archive Log Destination option is selected, IBM® Storage Defender Copy Data Management automatically manages the retention of only those archived logs that are under the new destination specified in the job definition. After a successful backup, logs older than the backup are automatically deleted from the IBM® Storage Defender Copy Data Management-managed destination.

If the Use Existing Archive Log Destination(s) option is selected in Oracle Backup job definition, IBM® Storage Defender Copy Data Management does not automatically purge any archived logs. The retention of archived logs must be managed externally, for example using RMAN. In order to support point-in-time recovery, ensure that the retention period is at least large enough to retain all archived logs between successive runs of the Oracle Backup job.

Oracle RMAN integration

Does the IBM® Storage Defender Copy Data Management Oracle solution integrate with RMAN?

Oracle Recovery Manager (RMAN), a command-line and Enterprise Manager-based tool, is the method preferred by Oracle DBAs for backup and recovery of Oracle databases, including maintaining an RMAN repository.

IBM® Storage Defender Copy Data Management creates application-consistent Oracle database copies simply – with no need to build and maintain complex RMAN protection scripts. At the same time, IBM® Storage Defender Copy Data Management automates cataloging of Oracle database copies in the RMAN recovery catalog. This enables DBAs to leverage RMAN for:

- Verification - IBM® Storage Defender Copy Data Management-created Oracle database copies can be instantly mounted so they can be easily verified through the RMAN verify command
- Advanced Recovery - IBM® Storage Defender Copy Data Management-created Oracle database copies can be instantly mounted to perform RMAN-driven PIT recoveries of database and tablespace.

IBM® Storage Defender Copy Data Management offers the following choice to the user for RMAN catalog registration:

- Register all copies in the RMAN catalog to enable RMAN recoveries against all copies
- Register on-demand only selected copies in the RMAN catalog to enable RMAN recoveries only when you need to recover

IBM® Storage Defender Copy Data Management simplifies full application-aware Oracle-consistent copy lifecycle while maintaining the flexibility and benefits of full RMAN-driven advanced recovery capabilities.

Pre and Post Scripts

Does IBM® Storage Defender Copy Data Management support pre/post scripts for the Application Backup workflow?

Yes, IBM® Storage Defender Copy Data Management supports job-level Pre/Post scripts and job-level pre/post Snapshot scripts to enable further customization.

- Job-level prescripts and postscripts are scripts that can be run before or after a job runs.
- Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot subpolicy runs. (Please refer to pre/post script topic in the IBM® Storage Defender Copy Data Management User's Guide for details.)

Data Masking

Does the IBM® Storage Defender Copy Data Management Oracle solution offer Data Masking integration with third party masking tools?

A concern for security officers in any organization is to keep confidential information locked down, even internally. Data masking is used to hide confidential data, by replacing it with fictitious data, when making data copies for DevTest or other use cases. It prevents leakage of sensitive data in non-production databases via static data masking [SDM], and production data in transit via dynamic data masking [DDM].

IBM® Storage Defender Copy Data Management includes integrated data masking workflows with the ability to leverage third party masking tools. Traditionally, data masking is difficult, slow, and storage-consuming, but with IBM® Storage Defender Copy Data Management it is easily integrated into the Oracle Backup workflow, allowing creation of masked copies at a specified frequency. Masked copies are automatically

marked in the Inventory. Access to secure copies is managed by the administrator by leveraging the application-level RBAC.

Is sample masking script provided with IBM® Storage Defender Copy Data Management?

A sample data masking script can be provided upon request. A sample masking script demonstrates data masking integration with built-in data masking functionality of Oracle 11g and 12c database system.

Oracle Restore Workflow

Can I leverage Oracle database clones for multiple use?

Limitations of current tools and approaches:

- Database cloning requires action by DBAs and is gated by process
 - QA relies on DBAs for cloning the databases for functional testing
 - The database cloning is gated by processes (space requisition, approvals, etc.)
- Database cloning is not time- or storage-efficient
 - Usage of RMAN or custom scripts creates full backup requiring large amounts of additional storage
 - Creating full copies is slow

IBM® Storage Defender Copy Data Management solves these challenges with a simple, automated, end-to-end clone lifecycle management:

- Self-service access to secure clones by QA team eliminates administrative and process bottlenecks
- IBM® Storage Defender Copy Data Management enables rapid database clones that are both time- and space-efficient
 - Provisions clones in minutes regardless of their size
 - Leverages underlying storage snapshots for space efficiency
- IBM® Storage Defender Copy Data Management promotes standardization and governance through centralized Inventory, granular RBAC, and automated jobs

Your Oracle clones can be utilized and consumed instantly – for whatever your use case -- through IBM® Storage Defender Copy Data Management “Instant Disk Restore” jobs. IBM® Storage Defender Copy Data Management catalogs and tracks all cloned instances. Instant Disk Restore can leverage iSCSI or FC protocol to provide an immediate mount of LUNs without transferring data.

Can I create an Instant Clone of an Oracle database for DevOps and Business Analytics?

IBM® Storage Defender Copy Data Management provides automated workflows to create instant clones of Oracle database regardless of its size.

- Instantly create database clones from any of the copies in the IBM® Storage Defender Copy Data Management Inventory, at local or remote locations, to accelerate Business Analytics.
- Enable and accelerate DevOps by providing Instant Disk Restore to secure clones of databases to appropriate users via application-level RBAC.

Then, when your TestDev, DevOps, or research/analytics work is completed, you can save the clone to more permanent storage or simply tear it down.

Can I create multiple clones from the same database snapshot?

Yes, multiple copies can be provisioned by IBM® Storage Defender Copy Data Management from the same snapshot. These copies can be either mounted with new DB names on the same Oracle Server or with same/different names on different Oracle servers.

How much space is used by my Oracle clones?

IBM® Storage Defender Copy Data Management utilizes zero-footprint instant clone capabilities provided by storage arrays. These clones do not use any additional space (except for maintaining pointers). Only new changes/writes to the read/writable clones utilize space.

What is the granularity of Database recovery supported by IBM® Storage Defender Copy Data Management?

Supported recoveries for standalone or RAC configurations

- Instant recovery of Oracle databases regardless of the size of the database

- Recover database(s) to original or to a new server (physical or virtual) simply with few clicks
- Recover database(s) with new names simply with few clicks
- Recover at DR site from replicated copies simply with few clicks
- Database can be recovered to point of snapshots to original or new location
- Perform PIT recoveries to original or new location simply with few clicks
- Recover RAC databases to any node on RAC simply with few clicks
- Database running on older versions can be recovered to an instance running same or newer version (IBM® Storage Defender Copy Data Management inherits any limitations defined by Oracle)
- Each selected database in a Restore job can have a separate destination specification. Databases can be recovered using a new name to an original or new instance
- You can select one or more databases in a single Restore job definition
- Recoveries are supported across the same storage type (i.e. ASM to ASM and standalone to standalone)
- Databases are always recovered in online mode

Users can additionally perform advanced fine-grained recoveries via RMAN integration. You can use IBM® Storage Defender Copy Data Management to instantly mount required snapshot copies to a specified Oracle server and use RMAN to recover. All RMAN-supported recoveries are available to users.

What granularity of Point-in-Time recovery is supported?

IBM® Storage Defender Copy Data Management enables you to recover databases to a specific point-in-time to enable you to:

- Restore to the state just before the point of failure.
- Restore multiple databases to a consistent time.

During PIT (point-in-time) Oracle database recoveries, if no log file snapshot exists that is newer than the chosen recovery PIT, the job creates a fresh log snapshot and uses it for recovery.

What happens to Database SID during recovery?

The Oracle System ID (**SID**) is used to uniquely identify a particular database on a server. One cannot have more than one database with the same SID on a single server. When using **RAC**, all instances belonging to the same database must have unique SIDs. A Restore job definition provides the user with an option to define a new name for a database during recovery. IBM® Storage Defender Copy Data Management will rename the SID if the user chooses a new database name during recovery, otherwise the original SID will be used.

How is the Oracle Initialization parameter file (init.ora) processed during Database recovery?

The Oracle initialization parameter file (init.ora) is created by the DBA and defines the overall instance configuration, such as how much memory should be allocated to the instance, the file locations, and internal optimization parameters.

IBM® Storage Defender Copy Data Management catalogs the initialization parameters during a job and uses them during recovery. IBM® Storage Defender Copy Data Management provides advanced options to control the behavior of processing initialization parameters used to start up the recovered database in Instant Database Recovery and DevOps workflows.

Customizable options allow you to use the same parameters as the source or specify a template pfile file to use.

Additionally, cluster-related parameters like instance_number, thread, and cluster_database are set automatically by IBM® Storage Defender Copy Data Management depending on the appropriate values for the destination.

For more information, see [“Renaming mount points and initialization parameter options” on page 250](#).

How does IBM® Storage Defender Copy Data Management set up mount points during Database recovery?

IBM® Storage Defender Copy Data Management provides advanced options for Instant Database Recovery and DevOps workflows to override the behavior of mount point handling during recovery. The Mount Point Rename option provides these selections:

- Append a timestamp: IBM® Storage Defender Copy Data Management appends a timestamp to the original mount point.
- Do not rename: IBM® Storage Defender Copy Data Management uses the same path/name for mount points or ASM diskgroups as the source.
- Add a custom prefix: Select this option to specify a custom prefix to be prepended to the source paths/names. The prefix value may contain leading or trailing slashes. In the case of ASM diskgroup names, the slashes are removed.
- Add a custom suffix: Select this option and specify a custom suffix to be appended the source paths/names.
- Replace a substring: Select this option to specify a custom string of characters in the old mount point to be replaced with another string of characters.

For examples of these options, see [“Renaming mount points and initialization parameter options” on page 250.](#)

Can I recover an Oracle Database running on a Linux® physical server to Oracle running on a Linux® VM?

Yes, the recovered database will be recovered as Physical RDM (pRDM).

Can I recover an Oracle Database running in Linux® VM to Oracle running on a Linux® Physical Server?

Yes, the pRDM configured database will be recovered as LUNs on the Physical Linux® Oracle server. Oracle configured with VMDK can only be recovered to VM.

Can I set up an Oracle Database to automatically be refreshed on a schedule?

The database "refresh" can be achieved by checking the "Allow overwrite and force cleanup of old session" option in the Restore job definition. You can run a job that spins up a database and the session will go into a Pending state awaiting cleanup. You can then run the same job again at a later time (either manually or scheduled) and it will automatically clean up the previous session before starting the new one. The cleanup process automatically stops and dismounts the database from the previous session. The new session will then mount and start the database again from the chosen snapshot (which is typically the latest as of runtime).

Where are Oracle specific and IBM® Storage Defender Copy Data Management specific logs if errors occur?

All required logs (IBM® Storage Defender Copy Data Management and Oracle application) are collected as part of the current log collection functionality. There should be no need to manually obtain Oracle application logs from within Oracle Servers.

Can a persistent agent be installed?

All required plugins are automatically injected on demand and automatically updated to the latest version if required.

Does Oracle application level encryption, such as Transparent Data Encryption(TDE), impact IBM® Storage Defender Copy Data Management?

Transparent Data Encryption (TDE) stops would-be attackers from bypassing the database and reading sensitive information from storage by enforcing data-at-rest encryption in the database layer. This is application-layer encryption that wouldn't typically impact IBM® Storage Defender Copy Data Management. The encryption keys live outside the database in a "wallet" that is managed separately by the administrator. IBM® Storage Defender Copy Data Management will create copies of the data on server A and mount them on server B, for example. The database software on server B should be able to read the encrypted data as long as the necessary keys are installed in the wallet there.

How do I refresh? How do I promote to Production?

All database recovery operations can leverage Instant Recovery mode (Test) and can then either be deleted or promoted to permanent mode via workflow control. This behavior is controlled via the job definition option Make Permanent.

- Enabled - Always make permanent
- Disabled - Never make permanent
- User election - Allows the user to select Make Permanent or Cleanup when the job session is pending

Does IBM® Storage Defender Copy Data Management provide Oracle specific reporting?

Reports are offered to assure that your Oracle databases are sound and that your IBM® Storage Defender Copy Data Management jobs are verified. Reports provided by IBM® Storage Defender Copy Data Management specifically for application support include:

- Application Configuration Report, which describes valuable system information about your Oracle Database Servers, and affirms that Oracle is configured correctly to be eligible for backup creation.

- Application RPO Compliance Report, which determines which of your Oracle database servers are not in compliance with your RPO parameters, and displays the reasons for their non-compliance.

System Requirements

What Oracle versions are supported and on what OS?

For Oracle versions supported on specific operating systems, see [Oracle requirements](#).

Does IBM® Storage Defender Copy Data Management support Oracle running on a VMware VM?

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

Does IBM® Storage Defender Copy Data Management support Oracle Database 12c Multitenant features?

IBM® Storage Defender Copy Data Management supports the Oracle Database 12c R1 Multitenant option for backup or clone of a container database (CDB). Recoveries of pluggable databases (PDB) are supported through RMAN.

How can granular recovery of specific PDBs be performed?

Granular recovery of specific Pluggable Databases (PDBs) can be performed via Instant Disk Restore recovery combined with RMAN. To do so:

1. Perform an Instant Disk Restore of the Container Database (CDB) by using an Oracle Restore job in IBM® Storage Defender Copy Data Management. The Oracle CDB backup may already have been cataloged into RMAN when the backup was created, or you can opt to do this during the Instant Disk Restore.
2. Login to RMAN. List the copies in the catalog and identify the tag from which you want to recover. Tags for IBM® Storage Defender Copy Data Management-created entries are generally of the form "ECX_<timestamp>".
3. Close the existing PDB by running the following command:

```
alter pluggable database <PDB_name> close;
```

4. Recover the PDB by running the following command:

```
run {
  restore pluggable database <name> from tag '<tag_name>';
  recover pluggable database <name>;
}
```

5. Open the recovered PDB by running the following command:

```
alter pluggable database <PDB_name> open;
```

Can I back up Oracle running on any storage to a supported storage system via VADP (VM Replication job)?

No, not through IBM® Storage Defender Copy Data Management Oracle Backup workflow. You can leverage a VMware Backup job with pre/post script to protect Oracle database in such configuration.

Oracle Requirements

Review requirements and pre-requisites for registering an Oracle provider. See [Oracle requirements](#)

Microsoft™ SQL Server Support FAQ

What is SQL Server CDM? How does it help solve my challenges?

Database administrators working extensively with SQL Server are challenged when faced with mission-critical use cases such as Backup, Recovery, DevOps, and Business Analytics. This is especially true given that their SQL Server databases have expanded in size and number over time, and that the databases need to be up and running 24x7x365.

SQL Server DBAs struggle with the following:

- Backups are slow, complex, and need constant management
- Backup process slows down the production servers
- Recoveries are slow and complex
- Repurposing App consistent backups (clones) for DevOps and Business Analytics is slow, complex and storage inefficient
- Lack of automation exists for producing quick and secure clones required to accelerate DevOps
- Copy sprawl problems occur due to no central catalog of copies
- Unable to meet organization's stringent RPO and RTO requirements

IBM® Storage Defender Copy Data Management simplifies SQL Server copy management by enabling administrators to orchestrate application-consistent copy creation, cloning and recovery in minutes, instead of hours or days. IBM® Storage Defender Copy Data Management copy management leverages the advanced snapshot and replication features of the underlying storage platform to rapidly create, replicate, clone, and restore copies of SQL Server databases in the most efficient way possible, in both time and space. IBM® Storage Defender Copy Data Management enables you to focus on the backup and restore requirements of your business rather than the technical details of the underlying storage platforms.

IBM® Storage Defender Copy Data Management is an intelligent copy data management solution that delivers end-to-end automation, orchestration, and self-service functionality for your SQL Server environment through a comprehensive and scalable catalog. With the self-service features of IBM® Storage Defender Copy Data Management, your users are empowered to create clones on demand, freeing DBAs, while at the same time offering the advanced recovery features needed for SQL Server environments.

IBM® Storage Defender Copy Data Management SQL Server Copy Data Management solution supports the following SQL Server deployment modes running on VMware virtual machines or physical servers:

- Standalone SQL Server – Databases running on a single server
- SQL Server Failover Cluster – SQL Server instances running on Windows™ Server Failover Clusters using Shared storage
- SQL Server Always On – Primary and secondary databases in Availability group configured across clusters of servers

Deployment and Registration

Do I need to deploy any additional agents to protect SQL Server standalone, Failover Cluster or AlwaysON configuration?

IBM® Storage Defender Copy Data Management for SQL Server is delivered as a VMware OVA that is easily deployed on demand in a matter of minutes. Once deployed, you simply register your SQL Servers with appropriate credentials and then let IBM® Storage Defender Copy Data Management discover the rest. IBM® Storage Defender Copy Data Management eliminates the complexity of manually deploying and maintaining application agents on SQL Servers. A lightweight application-aware agent is automatically injected and updated to the required SQL Servers on demand.

SQL Server Backup workflow

Application-consistent SQL Server backup (local and remote) - Step by step

IBM® Storage Defender Copy Data Management auto-discovers databases and enables copies only of eligible databases. To be eligible for IBM® Storage Defender Copy Data Management backup, the SQL Server database needs to be residing on a supported storage platform. With IBM® Storage Defender Copy Data Management, application owners do not need to be concerned about storage infrastructure.

A typical SQL Server database backup creation workflow consists of following steps:

- Auto-inject lightweight agent into SQL Server node running database instance
- Discover storage volume mapping to selected SQL Server database(s) and logs
- Place SQL Server database in hot backup mode via VMware Snapshot/VSS Snapshot (App consistent)
- Automatically create consistency group for related storage volumes (only on physical servers)
- Create application-consistent backup (VADP backup and/or Storage snapshot)
- Take SQL Server database out of hot backup mode (Delete VM snapshot/VSS snapshot)
- Optionally create log copies with lowest RPO possible
- Catalog SQL Server database backups in catalog
- Optionally replicate application-consistent copy to remote location leveraging storage replication feature

IBM® Storage Defender Copy Data Management creates and uses in-place copies, so no data is physically moved. IBM® Storage Defender Copy Data Management generated application-consistent copies are both space and time efficient. With the same ease, a DBA can automate the creation of remote copies for disaster recovery use cases.

Does SQL Server solution leverage the storage consistency group feature?

The storage consistency group feature allows storage administrators to take a snapshot of database applications where the data is spread across multiple volumes to maintain consistency across all volumes.

In a typical SQL Server Database, the data is spread across different volumes for better IO performance and availability. On Physical servers, IBM® Storage Defender Copy Data Management SQL Server application-consistent copy creation ensures that appropriate consistency groups are automatically created to maintain consistency across all related volumes. IBM® Storage Defender Copy Data Management SQL Server backup on VM relies on VMware snapshots and doesn't need to leverage storage consistency group feature.

What level of Application selection granularity is supported for SQL Server Backup jobs?

IBM® Storage Defender Copy Data Management SQL Server backup job definition supports copy selection at the following levels:

- One or more SQL Server Instances for Standalone SQL Server/Failover Cluster
- One or more Availability groups for AlwaysON
- One or more Databases for Standalone/Failover Cluster and SQL Server AlwaysON

Can I restore a database to an original instance and overwrite existing database in a single step?

Yes, use the Overwrite existing database option in the Application restore job definition.

Will IBM® Storage Defender Copy Data Management auto discover newly added SQL Server instances in a Standalone SQL Server and automatically protect it?

No. IBM® Storage Defender Copy Data Management will auto discover and present newly added SQL Server instances in the Backup job but you must explicitly select newly added SQL server instances for protection. Discovery of new SQL Server instances occurs as part of a regularly scheduled Application inventory job.

Will IBM® Storage Defender Copy Data Management auto discover newly added databases and automatically protect it?

Yes, if you select at Availability group level protection, IBM® Storage Defender Copy Data Management will auto discover newly added databases in selected availability group and protect it automatically during next job instance run. Discovery of new SQL Server instances and database occurs as part of regularly scheduled Application Inventory job.

Does IBM® Storage Defender Copy Data Management backup primary databases or secondary databases in SQL AlwaysOn?

IBM® Storage Defender Copy Data Management backs up only primary databases across the SQL AlwaysOn cluster.

Do SQL Server databases and logs need to be on supported storage for IBM® Storage Defender Copy Data Management CDM?

IBM® Storage Defender Copy Data Management also supports protection of SQL Server running on VMware VM configured on any storage that can be protected to supported storage systems via VM Replication. SQL Server running on physical servers require the database and logs to be on supported storage.

Does IBM® Storage Defender Copy Data Management perform full backups of databases?

IBM® Storage Defender Copy Data Management backups of SQL Server databases are always VSS COPY type backups.

SQL Server log management

Does IBM® Storage Defender Copy Data Management support Transaction log backup and log management?

Every SQL Server database has a transaction log that records all transactions and the database modifications made by each transaction. The transaction log must be truncated on a regular basis to keep it from filling up.

IBM® Storage Defender Copy Data Management provides you with an option to back up transaction log files. IBM® Storage Defender Copy Data Management supports log backup at a specified frequency. You can select one or more databases for log backup in a single backup job definition. Log destination can be specified as a single universal mount point or separate destination mount point for each database. Specified log backup destination path must already exist and must reside on supported storage system. If multiple databases are selected for backup, then each of the servers hosting the database must have their Destination directory set individually.

Does IBM® Storage Defender Copy Data Management support truncation of database logs?

Yes, IBM® Storage Defender Copy Data Management will automatically truncate log post log backups of databases that it backs up. If database logs are not backed up with IBM® Storage Defender Copy Data Management, its logs are not truncated by IBM® Storage Defender Copy Data Management and must be managed separately.

Can I specify a retention period for backed up transaction logs?

No. Log backup retention and auto-deletion is planned for a future release.

I am backing up transaction logs with IBM® Storage Defender Copy Data Management, but I don't want it to truncate logs. Can I control this behavior?

No. This will be enhanced in future release of IBM® Storage Defender Copy Data Management.

Pre and Post Scripts

Does IBM® Storage Defender Copy Data Management support pre/post scripts for Application Database Backup jobs?

Yes, IBM® Storage Defender Copy Data Management supports job-level pre/post scripts and job-level pre/post Snapshot scripts to enable further customization.

Job-level prescripts and postscripts are scripts that can be run before or after a job runs.

Snapshot prescripts and postscripts are scripts that can be run before or after a storage-based snapshot subpolicy runs. (Please refer to pre/post script topic in the IBM® Storage Defender Copy Data Management User's Guide for details.)

Data masking

Does IBM® Storage Defender Copy Data Management SQL Server solution offer Data Masking integration with third party masking tools?

A concern for security officers in any organization is that of keeping confidential information locked down, even internally. Data masking is used to hide confidential data, by replacing it with fictitious data, when making data copies for DevTest or other use cases. It prevents leakage of sensitive data in non-production databases via static data masking [SDM], and production data in transit via dynamic data masking [DDM].

The following Data Masking integration features will be available in a future release.

IBM® Storage Defender Copy Data Management will include integrated data masking workflows with the ability to leverage third party masking tools. Traditionally, data masking is difficult, slow, and storage-consuming, but with IBM® Storage Defender Copy Data Management it will be easily integrated into the SQL Server backup workflow, allowing creation of masked copies at a specified frequency. Masked copies are automatically marked in the catalog. Access to secure copies is managed by the administrator by leveraging the application-level RBAC.

In addition, SQL Server will enable you to leverage the Dynamic Data Masking feature of SQL Server 2016.

Is sample masking script provided with IBM® Storage Defender Copy Data Management SQL Server solution?

A sample data masking script can be provided upon request. A sample masking script demonstrates data masking integration with built-in Dynamic Data masking script of SQL Server 2016. This feature will be available in a future release.

SQL Server Restore Workflow

Can I leverage SQL Server database clones for multiple use?

Limitations of current tools and approaches:

- Database cloning requires action by DBAs and is gated by process
- QA relies on DBAs for cloning the databases for functional testing
- The database cloning is gated by processes (space requisition, approvals, etc.)
- Database cloning is not time- or storage-efficient
- Usage of common cloning tools or custom scripts creates full copy requiring large amounts of additional storage
- Creating full copies is slow

IBM® Storage Defender Copy Data Management solves these challenges with simple, automated end-to-end clone lifecycle management:

- Self-service access to secure clones by QA team eliminates administrative and process bottlenecks
- IBM® Storage Defender Copy Data Management enables rapid database clones that are both time- and space-efficient
- Provision clones in minutes regardless of its size
- Leverages underlying storage snapshots for space efficiency
- IBM® Storage Defender Copy Data Management promotes standardization and governance through centralized catalog, granular RBAC, and automated policies

Your SQL Server clones can be utilized and consumed instantly – for whatever your use case -- through IBM® Storage Defender Copy Data Management “Instant Disk Restore” jobs. IBM® Storage Defender Copy Data Management catalogs and tracks all cloned instances. Instant Disk Restore can leverage iSCSI or FC protocol to provide immediate mount of LUNs without transferring data.

Can I create an Instant Clone of a SQL Server database for DevOps and Business Analytics?

IBM® Storage Defender Copy Data Management provides automated workflows to create instant clones of SQL Server database regardless of its size.

- Instantly create database clones from any of the copies in the IBM® Storage Defender Copy Data Management inventory, at local or remote locations, to accelerate Business Analytics.
- Enable and accelerate DevOps by providing Instant Disk Restore to secure clones of databases to appropriate users via application-level RBAC.

Then, when your TestDev, DevOps, or research/analytics work is completed, you can save the clone to more permanent storage or simply tear it down.

What is the granularity of database recovery supported by IBM® Storage Defender Copy Data Management?

Supported recoveries for standalone or AlwaysOn:

- Database can be recovered to point of snapshots to original or new instance (Instant Disk Restore)

- Database can be recovered to point in time leveraging backed up transaction logs (Instant recovery) to original or new instance
- Database can be recovered using new name to original or new instance
- You can select one or more databases in a single restore job definition.
- Each selected database in a Restore job definition can have separate destination specification
- Databases are always recovered in online mode
- Database can be recovered from standalone instance to AlwaysON Availability group
- Database from AlwaysON Availability can be recovered to standalone instance
- Database running on older version can be recovered to instance running same or newer version.

What granularity of Point in Time recovery is supported?

IBM® Storage Defender Copy Data Management enables database recovery to a specific point in time, allowing you to:

- Restore to the state just before the point of failure
- Restore multiple databases to a consistent time

Can I restore a Database to a Transaction Mark?

No. This will be enhanced in a future release of the product.

Does IBM® Storage Defender Copy Data Management support recovering a database in online mode?

Yes, Instant Disk Restore or Instant Database Restore recovers databases in online mode.

Does IBM® Storage Defender Copy Data Management support recovering databases in an offline state (norecovery)?

No. This will be enhanced in a future release of the product.

Does IBM® Storage Defender Copy Data Management support recovering databases in a standby/read-only state (standby)?

Yes, IBM® Storage Defender Copy Data Management provides an application option to control this behavior. Roll back uncommitted transactions and leave the database ready to use

- Select this option to restore the database to an online state. If selected, additional transaction logs cannot be restored. If deselected, uncommitted transactions are not rolled back, leaving the database non-operational. Additional transaction logs can then be restored.

Does IBM® Storage Defender Copy Data Management support recovering database with Restricted Access?

No. This will be enhanced in a future release of the product.

Does IBM® Storage Defender Copy Data Management support restoring only logs so that they can be applied to a standby database?

No, not from the IBM® Storage Defender Copy Data Management application restore workflow. A user can easily access the transaction log backup location from the SQL server and perform this outside of the product.

Does IBM® Storage Defender Copy Data Management support out-of-place restore?

Out-of-place restore is used to relocate a database file to a new location:

- Copying/moving a database to a different location on a same SQL Server instance
- Copying a Database to a different SQL Server Instance at a different location

This feature will be enhanced in a future release of the product.

Where are SQL specific and IBM® Storage Defender Copy Data Management specific logs if errors occur?

All required logs (IBM® Storage Defender Copy Data Management and application) are collected as part of the current log collection functionality. There should be no need to manually obtain SQL application logs from within SQL Server VMs.

How do I refresh? How to promote to Production?

All database recovery operations can leverage Instant mode (Test) and then can either be deleted or promoted to permanent mode via workflow control. This behavior can be controlled via the Make Permanent job option.

- Enabled - Always make permanent
- Disabled - Never make permanent
- User election - Allows the user to select Make Permanent or Cleanup when the job session is pending

Does IBM® Storage Defender Copy Data Management use existing hardware providers for physical SQL backups?

No. IBM® Storage Defender Copy Data Management automatically deploys its own VSS HW provider service for SQL Server running on physical servers. It is automatically started on demand during SQL Server Backup jobs. At the completion of the backup job, the VSS HW provider service is automatically stopped.

When IBM® Storage Defender Copy Data Management protects a SQL VM with pRDM, can it restore a database back to the original node as a pRDM?

Currently, a SQL VM with pRDM must be registered as Physical in IBM® Storage Defender Copy Data Management. Hence, the restoration of that data obeys the Physical restore restrictions, which means it can only restore back to the original host via iSCSI. If the target host being restored to was registered as Virtual, then the database would be restored as a pRDM. This functionality will be improved in future release.

Why must I choose a proxy node when performing a restore to a SQL Failover cluster?

Windows™ requires signatures to be unique, so when you attach a disk that has a signature equal to one that is already attached, Windows™ keeps the disk in “offline” mode and doesn’t read its partition table or mount its volumes. To prevent disk signature collision, during Instant Database Restore, IBM® Storage Defender Copy Data Management leverages Windows™ proxy servers to temporarily mount disks from snapshots, generate a new signature, then mount to original server.

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

Self Service

Does IBM® Storage Defender Copy Data Management support RBAC? What is the level of granularity supported for SQL Servers?

Role-based access control allows you to set the resources and permissions available to IBM® Storage Defender Copy Data Management accounts. Through role-based access control you can tailor IBM® Storage Defender Copy Data Management for individual users, giving them access to the features and providers they need.

Using IBM® Storage Defender Copy Data Management RBAC functionality, user can delegate IBM® Storage Defender Copy Data Management role to enable and accelerate DevOps by providing Instant Access to secure clones of databases to appropriate users via application-level RBAC. Then, when your TestDev, DevOps, or research/analytics work is completed, you can save the clone to more permanent storage or simply tear it down.

Can developers access IBM® Storage Defender Copy Data Management operations using command line or APIs?

A rich set of REST APIs are provided to enable full access to IBM® Storage Defender Copy Data Management functionalities for further customization.

System Requirements

What SQL Server versions are supported and on what Windows™ OS? What are supported Storage Systems for Microsoft™ SQL Server?

For the most recent requirements and support, see [IBM Docs](#) for the version of IBM® Storage Defender Copy Data Management that is deployed.

What are Environment and permission requirements for SQL Server solution?

Note the following Microsoft™ environmental requirements:

- Windows™ Remote Shell (WinRM) must be enabled

- The SQL user must enable the public and sysadmin SQL permissions.
- The user identity must have sufficient rights to install and start the IBM® Storage Defender Copy Data Management Tools Service on the virtual machine node. This includes "Log on as a service" rights. For more information about the "Log on as a service" right, see <https://technet.microsoft.com/en-us/library/cc794944.aspx>.
- The fully qualified domain name must be resolvable and route-able from the IBM® Storage Defender Copy Data Management appliance
- The virtual machine node DNS name must be resolvable and route-able from the IBM® Storage Defender Copy Data Management appliance
- The VMGuest version must be current
- VMware Tools must be installed on the virtual machine node

Does IBM® Storage Defender Copy Data Management support SQL Server 2016 running on Windows™ 2016?

Yes. See the matrix above.

Does IBM® Storage Defender Copy Data Management support SQL Server configured as Physical RDMs, or Independent disks?

Yes. See footnotes 4 and 6 in the matrix above.

Does IBM® Storage Defender Copy Data Management support SQL Server configured as Virtual RDMs?

Yes. For limitations see footnote 5 in the matrix above.

Does IBM® Storage Defender Copy Data Management support SQL Server running on physical machine(s)?

Yes. See the matrix above.

Are there additional requirements for SQL support in IBM® Storage Defender Copy Data Management?

SQL Support for VMware Virtual Machines

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

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

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

In-Memory OLTP Requirements and Limitations

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

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

SQL Server Failover Clustering Requirements for Windows™ Server 2008 R2

The Failover Cluster Manager Snap-In must be imported and configured before running IBM® Storage Defender Copy Data Management Backup and Restore jobs. To import, run Windows™ PowerShell in Windows™ Server 2008 R2 and enter the following command: `import-module failoverclusters`. For more information, see [Microsoft.FailoverClusters.PowerShell](#).

Registration and Authentication

Register each SQL server as a provider in IBM® Storage Defender Copy Data Management by name or IP address. When registering a SQL Cluster (AlwaysOn), register each node by name or IP address. The fully qualified domain name and virtual machine node DNS name must be resolvable and route-able from the IBM® Storage Defender Copy Data Management appliance.

The user identity must have sufficient rights to install and start the IBM® Storage Defender Copy Data Management Tools Service on the node. This includes "Log on as a service" rights. For more information about the "Log on as a service" right, see [Add the Log on as a service Right to an Account](#).

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

Kerberos Requirements

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

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

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

Privileges

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

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

SAP HANA HSR Cluster Server Support FAQ

What is SAP HANA System Replication (HSR Cluster)?

SAP HANA System Replication (HSR) is a high-availability solution recommended by SAP to reduce downtime during planned maintenance, faults, or disasters. It ensures business continuity by replicating data between SAP HANA systems.

Key configuration details:

- SAP HANA HSR is set up as an active-passive cluster by using the Red Hat High Availability Add-On with Pacemaker.
- The primary node automatically owns the cluster's virtual IP address.
- Use the virtual IP address or DNS name to register the cluster with IBM® Storage Defender Copy Data Management.
- Additionally, register the IP addresses or DNS names of all participating cluster nodes. This is required for Instant Disk Restore (IDR) and in scenarios where the virtual IP becomes unavailable.
- The virtual IP or name is used across all workflows, including inventory, backup, and Instant Database Restore (IDBR) jobs.

What steps should be taken in IBM® Storage Defender Copy Data Management after a primary node failover in the SAP HANA HSR Cluster?

When a failover occurs:

- The secondary server becomes the new primary server.
- Pacemaker automatically updates the SAP HANA resources to reflect the new primary.

You must complete the following steps in the IBM® Storage Defender Copy Data Management UI:

1. Re-run inventory using the same virtual IP or name for the SAP HANA HSR cluster.
2. For restore operations, suspend HANA replication to prevent automatic failback. Then follow the appropriate restore type:
 - **Instant Database Restore (IDBR):** Suspend pacemaker before initiating the restore as IBM® Storage Defender Copy Data Management will attempt to shut down the running database. Ensure no automatic failover is triggered.
 - **Instant Disk Restore (IDR):** Suspend pacemaker before initiating the restore. Ensure no automatic failover is triggered at pacemaker.

What are the prerequisites for Instant Disk Restore (IDR)?

Before starting Instant Disk Restore (IDR) in IBM® Storage Defender Copy Data Management, complete the following steps:

On SAP Cluster:

1. Shut down HDB and dependent services on the secondary/passive node.
2. Stop cluster services:

```
systemctl stop pcsd
```

```
systemctl stop pacemaker
```

```
systemctl stop corosync
```

This will also stop sbd and hdb services.

3. Repeat the same steps on the primary/active node to stop the HDB and dependant services.
4. On IBM® Storage Defender Copy Data Management, Start the IDR job with default options.
5. After restore, start services on the primary/active node.:

```
systemctl start pcsd
```

```
systemctl start pacemaker
```

```
systemctl start corosync
```

This will automatically start **sbd** and **hdb** services.

6. Repeat the same steps on the **secondary/passive node**.
7. Once done, complete the replication configuration on both the nodes.

What are the prerequisites for Instant Database Restore (IDBR)?

Before attempting IBM® Storage Defender Copy Data Management Instant Database Restore (IDBR), complete the following instructions:

On SAP Cluster:

1. Shut down HDB and dependent services on the secondary node.
2. Stop cluster services:

```
systemctl stop pcsd
```

```
systemctl stop pacemaker
```

```
systemctl stop corosync
```

This will also stop sbd and hdb services.

3. On IBM® Storage Defender Copy Data Management, start the IDBR job with default options. Post Restore, IBM® Storage Defender Copy Data Management will start SAP services and make the database available for user connections. After restore operation, database should be up & running. Start services on the passive node:

```
systemctl start pcsd
```

```
systemctl start pacemaker
```

```
systemctl start corosync
```

4. Once done, complete the replication configuration on both the nodes.

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 Defender Copy Data Management family of products.

See the IBM® Storage Defender Copy Data Management [glossary](#).

Index

B

Backup and Restore the Catalog [350](#)

C

Creating PBHA SLA policy [96](#)

Creating PBR SLA policy [98](#)

L

Limitations and known issues [11](#)

T

Troubleshooting PBHA and PBR errors [341](#)

W

What's new 2.2.28 [11](#)

© Copyright International Business Machines Corporation 2017, 2025

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

