

PowerProtect Data Manager 19.14

Administrator Guide

Notes, cautions, and warnings

NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

Contents

Preface	8
Chapter 1: Getting Started	13
Introducing the PowerProtect Data Manager software.....	13
Supported Internet Protocol versions.....	14
Unsupported file-system modifications.....	17
References.....	17
Terminology.....	17
Access the PowerProtect Data Manager UI.....	18
Get Started window.....	19
UI tools and options	20
Dashboard	22
Export data.....	24
Exported fields.....	25
Customer feedback.....	26
Provide general feedback.....	26
Security configuration.....	26
Role-based security.....	27
Chapter 2: System Maintenance	28
Deploying and maintaining the health of PowerProtect Data Manager.....	28
Deploying and updating PowerProtect Data Manager.....	28
Licensing PowerProtect Data Manager.....	29
License types.....	29
Add a license.....	29
Specifying the PowerProtect Data Manager host.....	30
Specify a vCenter server as the PowerProtect Data Manager host.....	30
Minimum privileges required for the vCenter server PowerProtect Data Manager host.....	31
Memory optimization.....	31
Adjust the virtual machine memory.....	32
Restricted mode.....	32
System support.....	33
Configuring SupportAssist for PowerProtect Data Manager.....	33
Telemetry Collector	35
CloudIQ reporting.....	36
Set up the email server.....	36
Add AutoSupport.....	37
Enabling automatic update package checks and downloads.....	37
Add a log bundle.....	37
Audit logging and monitoring system activity.....	38
Monitor system services and system health.....	39
Access the open source software package information.....	44
Security certificates.....	44
Restarting PowerProtect Data Manager.....	45

System maintenance troubleshooting.....	45
Messages Catalog.....	45
Chapter 3: Managing Storage.....	46
Protection storage.....	46
PowerProtect DD Management Center automatic discovery.....	46
High Availability PowerProtect DD support.....	47
Smart Scale system pools.....	47
Add protection storage.....	50
Edit protection storage.....	51
Replace protection storage.....	51
Storage units.....	53
Storage unit limitations.....	54
Storage unit considerations for PowerProtect DD.....	54
Retention locking.....	54
Create a storage unit.....	55
Edit a storage unit.....	56
Delete a storage unit.....	57
Enable Indefinite Retention Hold on a storage unit.....	57
Disable Indefinite Retention Hold on a Storage Unit.....	58
Working with storage unit passwords.....	58
Differences in storage system and storage unit space reporting.....	58
Monitoring storage capacity thresholds.....	59
Chapter 4: Using the PowerProtect Search Engine.....	60
PowerProtect Search Engine.....	60
Set up and manage indexing.....	60
Search Engine node deletion.....	62
Delete an operational Search Engine node.....	62
Redeploy or delete a Search Engine node that failed to deploy.....	63
Delete all Search Engine nodes to remove the Search Engine cluster.....	63
Edit the network configuration for a Search Engine node.....	63
Perform a search.....	64
Troubleshooting Search Engine issues.....	64
Chapter 5: Managing Assets.....	70
About asset sources, assets, and storage.....	70
About other asset sources.....	71
Prerequisites for discovering asset sources.....	72
Discovering asset sources in an opaque network.....	72
Discovering asset sources in a GCVE environment.....	73
Full discovery of application asset sources.....	73
Enable an asset source.....	73
Disable an asset source.....	74
Delete an asset source.....	74
Adding a Cloud Snapshot Manager tenant.....	75
Add a Cloud Snapshot Manager tenant.....	75
Chapter 6: Managing Protection Policies.....	76

Protection policies.....	76
Before you create a protection policy.....	77
Replication triggers.....	80
Adding or editing a protection policy.....	81
Modify a policy name and description, objectives, or options.....	82
Changing storage targets.....	83
Redeploying storage targets.....	84
Replication to shared protection storage.....	84
Add or remove assets in a protection policy.....	85
Edit the retention period for backup copies.....	86
Viewing a summary of protection policies.....	86
View assets assigned to a protection policy.....	87
View the status of the last-run job of a protection policy.....	87
Run an asset-protection report.....	87
Add a service-level agreement.....	88
Run a compliance report.....	90
Disable a protection policy.....	91
Protection jobs running for a disabled policy.....	92
Enable a disabled protection policy.....	92
Customize the default behavior of disabled policies.....	92
Delete a protection policy.....	93
Overview of PowerProtect Data Manager Cloud Tier.....	93
Add a Cloud Tier objective to a protection policy.....	93
Manage Cloud Tier asset copies.....	94
Restore Cloud Tier backups to protection storage.....	95
Recall and restore from Cloud Tier.....	95
Extended retention for protection policies created in PowerProtect Data Manager 19.11 or earlier.....	95
Manual backups of protected assets.....	98
Manual backups of a single protected asset.....	98
Manual replication of protected assets.....	99
Manual Cloud Tiering of protected assets.....	100
Delete backup copies.....	100
Retry a failed backup copy deletion.....	101
Export data for deleted Oracle, SAP HANA and Storage Direct backup copies.....	102
Remove Exchange, File System, Kubernetes, Block Volume, and SQL backup copies from the PowerProtect Data Manager database.....	102
Removing expired backup copies.....	103
Removing assets from PowerProtect Data Manager.....	103
Remove assets and associated protection copies.....	104
Protecting client assets after a client hostname change.....	104
iGroup configuration and PowerProtect Data Manager policies.....	104
Troubleshooting failed replication jobs.....	106
Chapter 7: Restoring Data and Assets.....	108
View backup copies available for restore.....	108
Restoring a protection policy.....	109
Restore the PowerProtect Data Manager server.....	109
Restore Cloud Tier backups to protection storage.....	110
Recall and restore from Cloud Tier.....	110

Chapter 8: Preparing for and Recovering From a Disaster	112
About server disaster recovery.....	112
Differences between server DR methods.....	113
System recovery for server DR.....	113
Server DR protection storage types.....	113
Automatic server DR.....	114
Prepare the DD system recovery target (NFS).....	114
Manually configure server DR backups.....	115
Record settings for server DR.....	116
Manage PowerProtect Data Manager server DR backups.....	117
Recover PowerProtect Data Manager from server DR backups.....	117
Change the IP address or hostname of a DD system.....	121
Troubleshooting NFS backup configuration issues.....	123
Troubleshoot recovery of PowerProtect Data Manager.....	123
Recover a failed PowerProtect Data Manager restore.....	123
Disable server DR backups.....	124
Quick recovery for server DR.....	124
Quick recovery prerequisites.....	127
Identifying a remote system.....	128
Add a remote system for quick recovery.....	128
Edit a remote system.....	129
Quick recovery remote view.....	129
Troubleshooting failed quick-recovery jobs.....	130
Overview of PowerProtect Data Manager Cloud Disaster Recovery.....	130
Chapter 9: Managing Alerts, Jobs, and Tasks	131
Configure Alert Notifications.....	131
View and manage alerts.....	132
View and manage Audit Logs.....	133
Monitoring jobs and tasks.....	133
Monitor and view jobs and assets.....	134
View details for protection jobs.....	136
View details for asset jobs.....	138
View details for system jobs and tasks.....	140
Filter, group, and sort jobs.....	141
Restart a job or task manually.....	143
Restart a job or task automatically.....	144
Resume misfire jobs after a PowerProtect Data Manager update.....	145
Cancel a job or task.....	146
Exporting logs.....	147
Export logs for jobs.....	147
Export logs for assets or tasks.....	148
Limitations for alerts, jobs, and tasks.....	148
Chapter 10: Modifying the System Settings	149
System settings.....	149
Modify the network settings.....	149
Synchronizing the time between PowerProtect Data Manager and other systems.....	150

Modify the user-interface time zone, system time zone, and NTP server.....	150
Encryption in-flight.....	151
Server monitoring with syslog.....	153
Additional system settings.....	154
Modifying the PowerProtect Data Manager virtual machine disk settings.....	155
Modify the data disk size.....	155
Modify the system disk size.....	156
Configure the DD system.....	156
Virtual networks (VLANs).....	157
Virtual network traffic types.....	158
Virtual network topologies.....	160
Supported scenarios.....	163
Virtual network prerequisites.....	164
Configuring virtual networks.....	164
Virtual network asset assignment.....	168
Syslog server disaster recovery	170
Troubleshooting the syslog connection.....	170
No messages are transmitted to the syslog server.....	170
Chapter 11: Managing Reports.....	171
PowerProtect Data Manager reporting.....	171
Port requirements.....	171
Server requirements.....	172
Unsupported reporting engine vCenter operations.....	172
Known issues with the reporting engine and Report Browser.....	172
Configure and deploy the reporting engine.....	173
Updating the reporting engine from version 19.10.....	174
Report Browser.....	174
Types of reporting information.....	178
Filtering and customizing reports.....	179
Deleting the reporting engine.....	179
Managing disaster recovery of the reporting engine.....	179
Recover the reporting engine from a DR backup.....	179
Chapter 12: Configuring and Managing the PowerProtect Agent Service	183
About the PowerProtect agent service.....	183
Start, stop, or obtain the status of the PowerProtect agent service.....	184
Register the PowerProtect agent service to a different server address.....	185
Recovering the PowerProtect agent service from a disaster.....	185
Restore the PowerProtect Data Manager agent service datastore.....	186
Troubleshooting agent registration	186
Glossary.....	188

As part of an effort to improve product lines, periodic revisions of software and hardware are released. Therefore, all versions of the software or hardware currently in use might not support some functions that are described in this document. The product release notes provide the most up-to-date information on product features.

If a product does not function correctly or does not function as described in this document, contact Customer Support.

NOTE: This document was accurate at publication time. To ensure that you are using the latest version of this document, go to the Customer Support website.

Product naming

Data Domain (DD) is now PowerProtect DD. References to Data Domain or Data Domain systems in this documentation, in the user interface, and elsewhere in the product include PowerProtect DD systems and older Data Domain systems.

Isilon is now PowerScale. References to Isilon, Isilon products, or Isilon appliances in this documentation, in the user interface, and elsewhere in the product include PowerScale products and appliances.

In many cases the user interface has not yet been updated to reflect these changes.

Language use

This document might contain language that is not consistent with Dell Technologies current guidelines. Dell Technologies plans to update the document over subsequent future releases to revise the language accordingly.

This document might contain language from third-party content that is not under Dell Technologies control and is not consistent with the current guidelines for Dell Technologies own content. When such third-party content is updated by the relevant third parties, this document will be revised accordingly.

Acronyms

The acronyms used in this document might not be familiar to everyone. Although most acronyms are defined on their first use, a definition is not always provided with later uses of the acronym. For a list of all acronyms and their definitions, see the glossary at the end of the document.

Website links

The website links used in this document were valid at publication time. If you find a broken link, provide feedback on the document, and a Dell Technologies employee will update the link in the next release as necessary.

Purpose

The *Dell PowerProtect Data Manager Administrator Guide* describes how to configure, use, and administer PowerProtect Data Manager software.

Audience

This document is intended for the host system administrator who is involved in managing, protecting, and reusing data across the enterprise by deploying PowerProtect Data Manager software.

Revision history

The following table presents the revision history of this document.

Table 1. Revision history

Revision	Date	Description
01	July 11, 2023	Initial release of this document for PowerProtect Data Manager version 19.14.

Compatibility information

Software compatibility information for the PowerProtect Data Manager software is provided by the E-Lab Navigator.

Related documentation

The following publications are available at Customer Support and provide additional information:

Table 2. Related documentation

Title	Content
<i>PowerProtect Data Manager Administrator Guide</i>	Describes how to configure, use and administer the software. This guide also includes disaster recovery procedures. Procedures specific to asset protection are provided in the individual user guides.
<i>PowerProtect Data Manager Deployment Guide</i>	Describes how to deploy and license the software.
<i>PowerProtect Data Manager Release Notes</i>	Contains information about new features, known limitations, environment, and system requirements for the software.
<i>PowerProtect Data Manager Security Configuration Guide</i>	Contains security information.
<i>PowerProtect Data Manager Amazon Web Services Deployment Guide</i>	Describes how to deploy the software to Amazon Web Services (AWS).
<i>PowerProtect Data Manager Azure Deployment Guide</i>	Describes how to deploy the software to Microsoft Azure.
<i>PowerProtect Data Manager Google Cloud Platform Deployment Guide</i>	Describes how to deploy the software to Google Cloud Platform (GCP).
<i>PowerProtect Data Manager Cloud Disaster Recovery Administration and User Guide</i>	Describes how to deploy Cloud Disaster Recovery (Cloud DR), protect virtual machines in the AWS or Azure cloud, and run recovery operations.
<i>PowerProtect Data Manager Cyber Recovery User Guide</i>	Describes how to install, update, patch, and uninstall the PowerProtect Cyber Recovery software.
<i>PowerProtect Data Manager File System User Guide</i>	Describes how to configure and use the software with the File System agent for file-system data protection.
<i>PowerProtect Data Manager Kubernetes User Guide</i>	Describes how to configure and use the software to back up and restore namespaces and PVCs in a Kubernetes cluster or Tanzu Kubernetes cluster.
<i>PowerProtect Data Manager Microsoft Exchange Server User Guide</i>	Describes how to configure and use the software to back up and restore the data in a Microsoft Exchange Server environment.
<i>PowerProtect Data Manager Microsoft SQL Server User Guide</i>	Describes how to configure and use the software to back up and restore the data in a Microsoft SQL Server environment.

Table 2. Related documentation (continued)

Title	Content
<i>PowerProtect Data Manager Oracle RMAN User Guide</i>	Describes how to configure and use the software to back up and restore the data in an Oracle Server environment.
<i>PowerProtect Data Manager SAP HANA User Guide</i>	Describes how to configure and use the software to back up and restore the data in an SAP HANA Server environment.
<i>PowerProtect Data Manager Storage Direct User Guide</i>	Describes how to configure and use the software with the Storage Direct agent to protect data on VMAX storage arrays through snapshot backup technology.
<i>PowerProtect Data Manager Network-Attached Storage User Guide</i>	Describes how to configure and use the software to protect and recover the data on network-attached storage (NAS) shares and appliances.
<i>PowerProtect Data Manager Virtual Machine User Guide</i>	Describes how to configure and use the software to back up and restore virtual machines and virtual machine disks (VMDKs) in a vCenter Server environment with VADP or the Transparent Snapshots Data Mover (TSDM).
<i>PowerProtect Data Manager Storage Array User Guide</i>	Describes how to configure and use the software to protect and restore data on PowerStore storage arrays.
<i>VMware Cloud Foundation Disaster Recovery With PowerProtect Data Manager</i>	Provides a detailed description of how to perform an end-to-end disaster recovery of a VMware Cloud Foundation (VCF) environment.
PowerProtect Data Manager Public REST API documentation	Contains the Dell Technologies APIs and includes tutorials to guide you in their use.
<i>vRealize Automation Data Protection Extension for Data Protection Systems Installation and Administration Guide</i>	Describes how to install, configure, and use the vRealize Data Protection Extension.

Typographical conventions

The following type style conventions are used in this document:

Table 3. Style conventions

Formatting	Description
Bold	Used for interface elements that a user specifically selects or clicks, for example, names of buttons, fields, tab names, and menu paths. Also used for the name of a dialog box, page, pane, screen area with title, table label, and window.
<i>Italic</i>	Used for full titles of publications that are referenced in text.
<code>Monospace</code>	Used for: <ul style="list-style-type: none"> • System code • System output, such as an error message or script • Pathnames, file names, file name extensions, prompts, and syntax • Commands and options
<code>Monospace italic</code>	Used for variables.
Monospace bold	Used for user input.
[]	Square brackets enclose optional values.
	Vertical line indicates alternate selections. The vertical line means or for the alternate selections.
{ }	Braces enclose content that the user must specify, such as x, y, or z.
...	Ellipses indicate non-essential information that is omitted from the example.

You can use the following resources to find more information about this product, obtain support, and provide feedback.

Where to find product documentation

To find the latest documentation, navigate to the PowerProtect Data Manager Info Hub or type `www.dell.com/ppdmdocs` in your browser, or scan the following QR code on your mobile device.



Where to get support

The Customer Support website provides access to product licensing, documentation, advisories, downloads, and how-to and troubleshooting information. The information can enable you to resolve a product issue before you contact Customer Support.

To access a product-specific page:

1. Go to the Customer Support website.
2. In the search box, type a product name, and then from the list that appears, select the product.

Support Library

The Support Library contains a knowledge base of applicable solutions that you can search for either by solution number (for example, KB000xxxxxx) or by keyword.

To search the Support Library:

1. Go to the Customer Support website.
2. On the **Support** tab, click **Support Library**.
3. In the search box, type either the solution number or keywords. Optionally, you can limit the search to specific products by typing a product name in the search box, and then selecting the product from the list that appears.

Live chat

To participate in a live interactive chat with a support agent:

1. Go to the Customer Support website.
2. On the **Support** tab, click **Contact Support**.
3. On the **Contact Information** page, click the relevant support, and then proceed.

Service requests

To obtain in-depth help from a support agent, submit a service request. To submit a service request:

1. Go to the Customer Support website.
2. On the **Support** tab, click **Service Requests**.

NOTE: To create a service request, you must have a valid support agreement. For details about either an account or obtaining a valid support agreement, contact a sales representative. To find the details of a service request, in the **Service Request Number** field, type the service request number, and then click the right arrow.

To review an open service request:

1. Go to the Customer Support website.
2. On the **Support** tab, click **Service Requests**.
3. On the **Service Requests** page, under **Manage Your Service Requests**, click **View All Dell Service Requests**.

Online communities

For peer contacts, conversations, and content on product support and solutions, go to the Community Network. Interactively engage with customers, partners, and certified professionals online.

How to provide feedback

Feedback helps to improve the accuracy, organization, and overall quality of publications. You can send feedback to DPADDocFeedback@dell.com.

Getting Started

Topics:

- Introducing the PowerProtect Data Manager software
- Supported Internet Protocol versions
- Unsupported file-system modifications
- References
- Terminology
- Access the PowerProtect Data Manager UI
- Export data
- Customer feedback
- Security configuration

Introducing the PowerProtect Data Manager software

PowerProtect Data Manager software is an enterprise solution that provides software-defined data protection, deduplication, operational agility, self-service, and IT governance.

PowerProtect Data Manager key features include the following:

Table 4. Key features

Software-defined data protection with integrated deduplication, replication, and reuse
Data backup and recovery self-service operations from native applications that are combined with central IT governance
Multicloud optimization with integrated Cloud Tiering
SaaS-based monitoring and reporting
Modern services-based architecture for ease of deployment, scaling, and updating

PowerProtect Data Manager integrates multiple data-protection products within the Data Protection portfolio to enable data protection as a service, providing the following benefits:

Table 5. Benefits

Enables data-protection teams to create data paths with provisioning, automation, and scheduling to embed protection engines into their data-protection infrastructure for high-performance backup and recovery
Enables backup administrators of large-scale environments to schedule backups for the following asset types from a central location on the PowerProtect Data Manager server: <ul style="list-style-type: none"> • VMware virtual machines • File systems • VMAX storage groups • Kubernetes clusters • Microsoft Exchange Server and Microsoft SQL Server databases • Oracle databases • SAP HANA databases • Network-attached storage (NAS) shares • PowerStore block volumes
Provides an agent-based approach to automatically discover and protect databases on an application server
Enables self-service and centralized protection by: <ul style="list-style-type: none"> • Monitoring service-level objectives (SLOs)

Table 5. Benefits (continued)

<ul style="list-style-type: none"> Identifying violations of recovery-point objectives (RPOs)
Supports deploying an external VM Direct appliance that moves data with a VM Direct Engine that is optimized for performing high-capacity backup streams
Comes with a basic embedded VM Direct Engine that has the following functions and capabilities: <ul style="list-style-type: none"> It is automatically used as a fallback proxy for performing backup and restore operations when an external VM Direct Engine fails, is disabled, or is unavailable It has a limited capacity for performing backup streams It can work with virtual-machine crash-consistent protection policies that use the Transparent Snapshot Data Mover (TSDM) protection mechanism It enables the Search Service used by PowerProtect Search
Supports PowerProtect Search, which enables backup administrators to quickly search for and restore VM and NAS file copies
Supports the vRealize Automation DP extension, which enables the automatic provisioning of virtual machines and on-demand backups and restores
Integrates with Cloud Disaster Recovery (Cloud DR), including workflows for Cloud DR deployment, protection, and recovery operations in the AWS and Azure clouds
Integrates with PowerProtect Cloud Snapshot Manager to view PowerProtect Cloud Snapshot Manager jobs, alerts, and reports from a consolidated PowerProtect Data Manager dashboard
Integrates with PowerProtect Cyber Recovery to protect the integrity of a PowerProtect Data Manager environment from cyber threats
Provides a RESTful API interface that allows PowerProtect Data Manager to be monitored, configured, and orchestrated: <ul style="list-style-type: none"> Existing automation frameworks can be integrated New scripts can be quickly written Easy-to-follow tutorials are provided

Supported Internet Protocol versions

PowerProtect Data Manager and its components support IPv4 and IPv6 addresses in certain configurations.

Table 6. Supported configurations

Component	Internet Protocol
PowerProtect Data Manager core	IPv4 only or both IPv4 and IPv6
PowerProtect Data Manager cloud deployments (AWS, Azure, GCP)	IPv4 only NOTE: Despite other entries in this chart to the contrary, if PowerProtect Data Manager is deployed to a cloud environment, no component in the cloud can use IPv6.
VM Direct, TSDM, and Search	IPv4 only or IPv6 only NOTE: Virtual machines that are backed up must use the same protocol that VM Direct uses. Virtual machines can use both IPv4 and IPv6, even though VM Direct and TSDM cannot.
Application agents integrated with PowerProtect Data Manager:	NOTE: If both IPv4 and IPv6 are configured and the PowerProtect Data Manager FQDN is used, the agent uses IPv6 for network communication.
<ul style="list-style-type: none"> File System 	IPv4, IPv6, or both
<ul style="list-style-type: none"> Microsoft Exchange Server 	IPv4 only or both IPv4 and IPv6
<ul style="list-style-type: none"> Microsoft SQL Server (Application Direct) 	IPv4, IPv6, or both
<ul style="list-style-type: none"> Microsoft SQL Server (VM Direct) 	IPv4 only or IPv6 only

Table 6. Supported configurations (continued)

Component	Internet Protocol
<ul style="list-style-type: none"> • Oracle RMAN • SAP HANA • Storage Direct 	<p>① NOTE: Only the Microsoft SQL Server agent supports VM Direct.</p> <p>IPv4, IPv6, or both</p> <p>IPv4, IPv6, or both</p> <p>IPv4 only</p>
Stand-alone application agents	IPv4 only
Network-attached storage (NAS)	IPv4, IPv6, or both
Storage arrays (PowerStore)	IPv4 only
Kubernetes	IPv4 only
PowerProtect Data Manager management	IPv4 or IPv6
PowerProtect DD communication	IPv4 or IPv6
Report Browser	<p>IPv4 only</p> <p>① NOTE: If PowerProtect Data Manager is configured to use both IPv4 and IPv6, configuring an NTP server and setting a time zone is required for accurate date and time information in reports.</p>
SupportAssist	IPv4, IPv6, or both
Syslog Log Server Gateway	IPv4 or IPv6

The following limitations and considerations apply.

Communication with components

If PowerProtect Data Manager is configured to only use one protocol, all components it communicates with must also use that protocol. If some components that PowerProtect Data Manager communicates with use IPv4 and others use IPv6, PowerProtect Data Manager must be configured to use both IPv4 and IPv6.

DD systems and DDVE

If a DD system or a DDVE instance uses only IPv6, the required IPv6 interface must be manually selected when a protection policy is added or edited.

Network-attached storage and DD-system storage units

If the storage unit of a protection policy is different or changed from the destination asset source, you must assign a network to the destination asset for a successful restore. For example, if your source asset is backed up in an IPv6 network, you must assign an IPv6 network to the destination asset for the restore to be successful.

To assign a network for the destination asset, perform the following steps:

1. In the PowerProtect Data Manager UI, select **Infrastructure > Assets > NAS**.
2. Select the destination asset, click **More Actions** and select **Assign Network**. The **Assign Network** page appears.
3. Select a network from the **Network Label** list, click **Save**.
4. If a restore failed because of the wrong destination address, retry the operation.

Disaster recovery

Recovering a PowerProtect Data Manager server might result in a conflict with protection-policy configurations. For instance, if the recovered server is configured to use only IPv4, a protection policy that is configured to use IPv6 cannot run.

Name resolution

Name resolution and reverse IP lookup must be configured to ensure the following:

- Fully qualified domain names of PowerProtect Data Manager, its components, and DD components resolve to a valid IPv4 or IPv6 address.
- If both IPv4 and IPv6 addresses are used for DD, both addresses resolve to the same FQDN.
- All IPv4 and IPv6 addresses are valid and reachable.
- The FQDNs of application-agent hosts that use FQDN as their preferred host address resolve to a valid IPv4 or IPv6 address.
- Each application-agent host that uses FQDN as its preferred host address resolves the FQDN of PowerProtect Data Manager to an IP address of the same protocol that it uses. For example, if a host uses IPv4, it resolves the FQDN of PowerProtect Data Manager to an IPv4 address.

Server updates

IPv6 is only supported with new deployments of PowerProtect Data Manager 19.12 or later. Using IPv6 after updating from PowerProtect Data Manager 19.11 or earlier is unsupported.

Search Engine indexing and adding IPv6 to an IPv4-only system

If you add IPv6 to an IPv4-only system, indexing from any existing Search Engine cluster becomes unavailable. After adding IPv6, you must delete all IPv4 Search Engine nodes to remove the Search Engine cluster, and then add new IPv6 nodes to a new cluster.

Unlike other PowerProtect Data Manager components, if IPv6 is used with a Search Engine, the FQDN of all Search Engine nodes and related DD systems must always resolve to an IPv6 address and never to an IPv4 address.

Storage Policy Based Management

If using vCenter or ESXi 7.0u2 or earlier with only IPv6, SPBM providers must be added using their PowerProtect Data Manager FQDN.

Service Unavailable messages with the vSphere Client PowerProtect plug-in

If vCenter uses the vSphere Client PowerProtect plug-in with IPv6 and the vCenter host is added to PowerProtect Data Manager using its IPv6 address or FQDN, *Service Unavailable* messages might be seen for the protected virtual machine. Backups and restores of the protected virtual machine are unaffected, and these messages can be ignored.

Uncompressed IPv6 formatting

Network interfaces that exist on a DD 7.4.x or earlier system and that are configured to use an uncompressed IPv6 format cannot be discovered. An example of an uncompressed IPv6 format is 2620:0000:0170:0597:0000:0000:0001:001a. An example of a compressed IPv6 format is 2620:0:170:597::1:1a. To use these network interfaces, reconfigure them to use either an IPv4 address or a compressed IPv6 address, and then initiate a discovery.

Unsupported file-system modifications

Files and directories on PowerProtect Data Manager and PowerProtect DD systems should only be modified according to documentation and guidance.

Performing any of the following file-system operations that have not been documented in a product guide or communicated by Customer Support is unsupported:

- Adding, removing, editing, or otherwise modifying a file or directory
- Manually mounting a DD file system with anything other than read-only permissions
- Altering a file-system procedure
- Replacing a command in the step of a file-system procedure with a different command

References

Some procedures in this document reference other publications for further details.

For a list of PowerProtect Data Manager publications, see "Related documentation" in the preface.

For information about DD Virtual Edition, see the following publications at Customer Support:

Table 7. Related PowerProtect DD Virtual Edition documentation

<i>PowerProtect DD Virtual Edition in VMware Cloud Installation and Administration Guide</i>
<i>PowerProtect DD Virtual Edition in Google Cloud Platform Installation and Administration Guide</i>
<i>PowerProtect DD Virtual Edition on Premise Installation and Administration Guide</i>
<i>PowerProtect DD Virtual Edition in Azure Installation and Administration Guide</i>
<i>PowerProtect DD Virtual Edition in Amazon Web Services Installation and Administration Guide</i>

Terminology

Familiarize yourself with the terminology for the PowerProtect Data Manager user interface and documentation.

The following table provides more information about names and terms that you should know to use PowerProtect Data Manager:

Table 8. Term list

Term	Description
Application agent	Application agents are installed on application or database host servers to manage protection using PowerProtect Data Manager. These agents are commonly known as DD Boost Enterprise Agents (DBEAs) for databases and applications.
Application-aware	A virtual machine protection policy that includes additional application-aware data protection for Microsoft SQL Servers. An application-aware virtual machine protection policy provides the ability to quiesce the application during virtual machine image backup to perform a full backup of Microsoft SQL Server databases. You can also schedule Microsoft SQL Server log backups for the virtual machines in the policy.
Asset	Assets are objects in PowerProtect Data Manager for which you want to manage protection, including virtual machines, databases, and file systems.
Asset source	Assets that PowerProtect Data Manager protects reside within asset sources, which include vCenter servers, application or database hosts, and file servers.
Cloud Tier storage	Cloud Tier storage can be added to a protection storage system to expand the deduplication storage capacity onto less expensive object storage in public or private object storage clouds, including secure Elastic Cloud Storage appliances.
Copy	A PowerProtect Data Manager copy is a point-in-time backup copy of an asset.

Table 8. Term list (continued)

Term	Description
Copy Map	The PowerProtect Data Manager Copy Map is a visual representation of backup copy locations on your protection storage and is available for all protected assets that have copies.
Discovery	Discovery is an internal process that scans asset sources to find new assets to protect and scans infrastructure components to monitor their health and status.
Instant Access	PowerProtect Data Manager virtual machine backup copies can be accessed, mounted, and booted directly from the protection storage targets as running virtual machines. This operation is called Instant Access. Copies can also be moved to a production VMware datastore using vMotion. PowerProtect Data Manager Virtual machine application-aware backup copies can be mounted directly from protection storage as running Microsoft SQL Server databases, which includes the ability to roll forward log backups. These Microsoft SQL Server database disks can also be moved to a production VMware datastore using vMotion.
PowerProtect Data Manager agent	An agent that is included in PowerProtect Data Manager and installed on each application agent host server so that you can monitor and manage the application agent through PowerProtect Data Manager.
Protection policy	Protection policies configure and manage the entire life cycle of backup data, which includes backup types, assets, backup start and stop times, backup devices, and backup retention.
Service-level agreement (SLA)	An optional policy that you can layer on top of a protection policy. An SLA performs additional checks on protection activities to ensure that protection goals meet the standards of an organization. SLAs are made up of one or more service-level objectives.
Service-level objective (SLO)	A definable rule that sets the criteria for recovery-point objectives (RPOs), encryption, and the location of backups according to company requirements.

Access the PowerProtect Data Manager UI

PowerProtect Data Manager provides a web-based UI that you can use to manage and monitor system features and settings from any location over a network.

About this task

Note that after 30 minutes of inactivity, this interface might fail to respond or you might see one of the following errors:

- 401: Authentication Required
- 503: Unknown Error

To resolve any of these issues, refresh your browser and log in. If you logged in before, you need to log in again.

Steps

1. From a host that has network access to the virtual appliance, use the latest version of Google Chrome to connect to the appliance:

`https://<appliance_hostname>`



NOTE: You can specify the hostname or the IP address of the appliance.

2. Log in with your username and password.

Username follows the format `user[@domain]`, where `domain` is an optional identifier that associates the user with a particular identity provider.

For example: `jsmith` or `administrator@test-lab`.

- If you do not supply a domain, the authentication service checks the default identity provider.
- If you supply a domain, the authentication service consults the external identity provider for that domain and determines whether to allow the login.
- If multi-factor authentication (MFA) is enabled, the **Multi-Factor Authentication** dialog box prompts you for a passcode. PowerProtect Data Manager verifies this passcode with the MFA service before allowing the login.



NOTE:

If you log in with an expired password, reset the password immediately. Clicking **Cancel**, closing the browser, or navigating away from the page before changing your password disables your credentials for subsequent logins. If you log in and receive a prompt to change your password because of outdated login credentials, provide your current password, a new password, and confirmation of the new password to continue.

When the identity provider validates the credentials, the authentication service issues a user token. The PowerProtect Data Manager UI uses the token information to authorize activities.

Unless you have changed the system configuration, the default identity provider is the local identity provider.

The *PowerProtect Data Manager Security Configuration Guide* provides more information about the available user roles and their associated permissions. The associated roles for an account determine what parts of the UI a user can see and use, and what operations a user can perform.

If this is your first time accessing the PowerProtect Data Manager UI, an unsigned certificate warning might appear in the web browser.

The security certificate that encrypts communication between the PowerProtect Data Manager UI and the web browser is self-signed. A self-signed certificate is signed by the web server that hosts the secure web page. There is nothing wrong with this certificate. This certificate is sufficient to establish an encrypted channel between the web browser and the server. However, it is not signed by a trusted authority.


The **Get Started** window appears with configuration options that are required on first deployment. To skip this window and go right to the **Dashboard**, click **Launch**.

From the **Dashboard** window:

- The left pane provides links to the available menu items. Expand a menu item for more options.
- The icons in the PowerProtect Data Manager banner provide additional options.

Get Started window

The **Get Started** window provides configuration options that are required when the PowerProtect Data Manager system is first deployed. This window continues to display by default each time you log in until you click **Launch**.

You can access the **Get Started** window at any time, or view any getting started options that have yet to be configured, by clicking , and then selecting **Getting Started**.

The **Get Started** window enables you to configure or edit the following menu items:

Table 9. PowerProtect Data Manager Get Started menu items

Options	Description
License	Launches the License window, which prompts you to add a license file to PowerProtect Data Manager. Once a license is uploaded, you can view license details, such as capacity usage and software ID.
Support	Launches the Support window, which enables you to configure SupportAssist, AutoSupport, and set up the email server for application notifications and messages.
Assets	Launches the Asset Sources window, where you can enable any of the asset source types that PowerProtect Data Manager supports. After enabling an asset source, you can add and register the source for the protection of assets.
Storage	Launches the Add Storage window, where you can add a PowerProtect DD System or PowerProtect DD Management Center as protection storage for primary backup and replicated copies.

UI tools and options

Learn about the tools, windows, and banner options available in the PowerProtect Data Manager UI.

PowerProtect Data Manager UI tools and windows

The following table describes the tools and windows in the PowerProtect Data Manager UI left navigation pane.

Table 10. PowerProtect Data Manager tools











Menu item	Description
 Dashboard	Click Dashboard to view the overall state of the PowerProtect Data Manager system.
 Health	Click Health to view a score for the overall PowerProtect Data Manager system health (Good, Fair, or Poor).
 Infrastructure	Click Infrastructure to: <ul style="list-style-type: none">View and manage all assets:<ul style="list-style-type: none">VMware virtual machinesFile systemsVMAX storage GroupsNOTE: Not applicable to the PowerProtect Data Manager appliance.Kubernetes clustersMicrosoft Exchange Server databasesNetwork Attached Storage (NAS)Microsoft SQL Server databasesOracle databasesSAP HANA databasesBlock volumesAdd vCenter and application and File System host asset sources.View and manage Integrated Storage.Add a VM Direct appliance with the VM Direct protection engine for virtual machine data protection.Manage the vSphere Installation Bundle (VIB) for virtual machine crash-consistent data protection performed with the Transparent Snapshot Data Mover (TSDM) protection mechanism.Manage registration of Oracle RMAN agent, Microsoft application agent, SAP HANA agent, and File System agent.View and manage Cloud Disaster Recovery.Create and manage a Search Cluster.Add PowerProtect Cloud Snapshot Manager tenants as asset sources for jobs, alerts, and reports.
 Protection	Click Protection to: <ul style="list-style-type: none">Add protection policies to back up assets.Manage service-level agreements (SLAs).Add, edit, and delete protection rules for asset inclusion in policies.Add, edit, and delete file exclusion templates for File System protection policies.
 Restore	Click Restore to: <ul style="list-style-type: none">View asset copy location details and initiate a Restore operation.Manage Instant Access Sessions.Use the File Search feature to find and restore virtual machine file copies.

Table 10. PowerProtect Data Manager tools (continued)

Menu Item	Description
 Alerts	Click Alerts to: <ul style="list-style-type: none"> • View and acknowledge alerts and events. • Filter alerts by critical, warning, and informational status, and specify the time range. • View and examine Audit logs. • Export audit logs to .csv files. • Set audit log boundaries. • Configure alert notifications. There is also a banner UI option, represented by  , which provides links that enable you to view all unacknowledged alerts.
 Administration	Click Administration to: <ul style="list-style-type: none"> • Configure users and roles. • Set password credentials and manage key chains. • View and replace certificates. • Add external identity providers. • View and manage resource groups.
 Reports	Click Reports to access the PowerProtect Data Manager Report Browser and Reporting Engine .
 Jobs	Click Jobs to manage jobs, view by protection or system, filter, and view details.

Banner UI options

The following table describes the icons in the PowerProtect Data Manager UI banner.

Table 11. Banner UI options









Option	Description
	Click to provide customer feedback.
	Click to enter search criteria to find assets, jobs, logs, and alerts.
	The number next to this icon indicates the critical unacknowledged alerts over the last 24 hours. Click to expand for more information about unacknowledged alerts, including: <ul style="list-style-type: none"> • The total number of alerts (all statuses — critical, warning, or informational) that have yet to be acknowledged, or just the unacknowledged alerts from the last 24 hours (marked with the New tag). • The number of critical alerts that have yet to be acknowledged, or just the unacknowledged critical alerts from the last 24 hours (marked with the New tag). Within this menu, click any of these links to open the Alerts window, where you can view specific details about these unacknowledged alerts.
	Click to restore assets from replicated copies through quick recovery. This icon only appears when the system receives replicated metadata from a source system.

Table 11. Banner UI options (continued)

Option	Description
	Click to configure and manage PowerProtect Data Manager system network, time zone, and NTP settings, DR backups, security, licenses, updates, authentication, agent downloads, and support, and to access the Get Started window.
	Click to obtain more information about PowerProtect Data Manager, access Customer Support, or view the REST API documentation.
	Click to log out, and log in as a different user, or change the current user password.
	Click to launch CloudIQ, APEX Backup Services, Cloud Snapshot Manager, or vProtect. NOTE: Only Cloud Snapshot Manager is applicable to the PowerProtect Data Manager appliance.

Dashboard

The Dashboard is visible when you log in to the PowerProtect Data Manager UI, and can be accessed from the left navigation pane.

The **Dashboard** window provides a high-level view of the overall state of the PowerProtect Data Manager system through six widgets. The following table describes each widget.

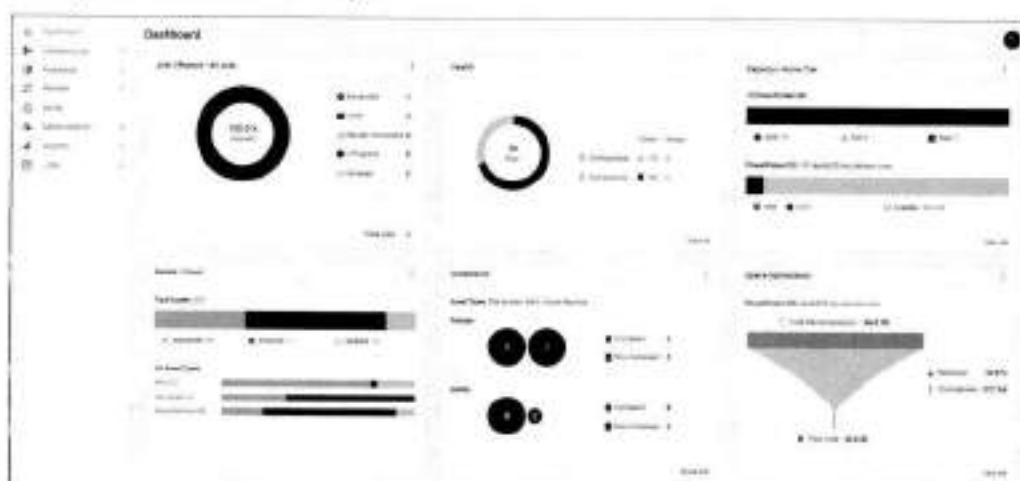


Figure 1. Dashboard widgets

Table 12. PowerProtect Data Manager Dashboard

Dashboard widget	Description
Jobs Protection Jobs Restore Jobs System Jobs Asset Level	<p>This widget provides a color-coded status of backup, restore, and system jobs that are in progress or have been performed in PowerProtect Data Manager over a specified period. Jobs Protection displays by default, showing jobs performed over the last 24 hours.</p> <p>Click the three vertical dots at the top of the widget to:</p> <ul style="list-style-type: none"> • Select Protection, Restore, System or Asset Level to switch the jobs view in the widget. • Choose the time period for the jobs that you want to view (last 24 hours, last 3 days, last 7 days, or all). Once a time period is selected, the widget updates to display only jobs performed within that time period. <p>Click a color in the chart to view details about jobs with a specific status, or click the links next to each status. This will open the appropriate Jobs window, which is filtered to display the jobs that match the selected status and time period. From this window, you can manage jobs, view more details, and search jobs.</p>

Table 12. PowerProtect Data Manager Dashboard (continued)


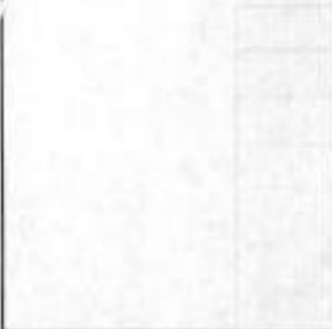


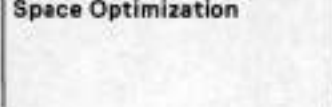
Dashboard widget	Description
<p>Assets Count and Assets Size</p> 	<p>Details in this widget include the number of protected assets, unprotected assets, and excluded assets for each asset source that has been added and enabled in PowerProtect Data Manager. You can also view the total number of assets for each asset source, and the total size of these assets. Assets Count displays by default, and the asset types are sorted based on the percentage of the total asset count that are unprotected, or the total size of the unprotected assets for the asset source, depending on the view.</p> <p>Click the three vertical dots at the top of the widget to:</p> <ul style="list-style-type: none"> • Select Count or Size to switch the assets view in the widget. • Select one or more asset sources from the list. You can display asset statistics for a single asset source, multiple asset sources, or all asset sources. <p>Hover over a color to view the exact number of protected, unprotected, and excluded assets and the total size of these assets. Click a color to open the Infrastructure > Assets window, which is filtered to display the assets that match the selected status.</p>
<p>Health</p> 	<p>This widget provides a score for the overall PowerProtect Data Manager system health (Good, Fair, or Poor). Health details and status are provided for the following categories:</p> <ul style="list-style-type: none"> • Components: Identifies the state of hardware and software services, such as <i>Running</i> or <i>Failed</i>. • Configuration: Identifies whether any aspects of the PowerProtect Data Manager configuration are incomplete, such as <i>System Support</i> configuration. • Capacity: Identifies the provisioned and currently allocated size of the associated storage system. • Performance: Identifies key performance indicators, such as <i>memory use</i>. • Data Protection: Identifies key protection indicators, such as <i>service-level agreements not being met</i> and <i>disaster-recovery backup copies not being present</i>. <p>Click View All to view more details about the system health issues for all categories.</p>
<p>Compliance</p> 	<p>This widget provides compliance verification statistics for protection policies that are linked to a Service Level Agreement (SLA). The widget also identifies the number of assets within these policies that are compliant and non-compliant.</p> <p>Click the three vertical dots at the top of the widget to select one or more asset sources from the list. You can display compliance statistics for a single asset source, multiple asset sources, or all asset sources. By default, the total count and number of protection policies for compliant and non-compliant assets displays for all asset sources.</p> <p>Click View All to open the Protection > SLA Compliance window, where you can view more details about the specific policies and assets that are non-compliant.</p>
<p>Capacity Active Tier and Capacity Cloud Tier</p> 	<p>This widget displays the capacity status of the DD protection storage systems that are associated with this instance of PowerProtect Data Manager for the active tier and cloud tier. Based on the available capacity on each DD system, a color coded bar graph displays the number of systems that are Good (>20% available), Fair (<20% available), or Poor (<10%).</p> <p>Click the three vertical dots at the top of the widget to:</p> <ul style="list-style-type: none"> • Select Active Tier or Cloud Tier to switch between a view of protection storage systems for the active tier and cloud tier in the widget. By default, the widget displays Capacity Active Tier. • Select a DD system from the list. The widget updates to display capacity statistics for the selected DD system. You can only display capacity statistics for one system at a time. <p>Click View All to open the Infrastructure > Storage window, where you can view more details about specific protection storage systems.</p>
<p>Space Optimization</p> 	<p>This widget provides information about how efficient the active tier storage capacity is on individual DD systems associated with this instance of PowerProtect Data Manager. Efficiency is determined based on the size of pre-compression data compared with the size of post-compression data on the system.</p>

Table 12. PowerProtect Data Manager Dashboard (continued)

Dashboard widget	Description
	Click the three vertical dots at the top of the widget to select a DD system from the list. The widget updates to display space optimization statistics for the selected DD system.

Export data

PowerProtect Data Manager enables you to export and save table data in CSV format.

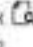
Prerequisites

In the PowerProtect Data Manager UI, browse to a window that includes the **Export All** functionality.

About this task

The following table lists the windows that support the **Export All** functionality.

Table 13. Supported windows

Menu item	Window
Health	Health
Infrastructure	Assets
	Application Agents
Protection	Protection Policies
	You can also export records for assets that are assigned to a protection policy. Select a protection policy to view its details, and then click the asset count link next to Assets .
	SLA Compliance
Protection	Protection Rules
	You can also export records for assets that are applied to a protection rule. Click the link in the Assigned Assets Count column for the protection rule.
Restore	Assets
Alerts	System
Administration	Access Control > Users/Groups
	Access Control > Resource Groups
	You can also export records for assets that are assigned to a resource group. Click  next to the resource group, and then click View Assets in the right pane.
Administration	Audit Logs
Jobs	Protection Jobs
	System Jobs
System Settings	Messages Catalog

Steps

- (Optional) Filter and sort the information that appears in the table.
- In the window, click **Export All** to export the data to a .csv file.

NOTE: Filters applied to the table in the **Protection Policy** window are not applied to the exported .csv file. Exported protection records include all data that is shown in the table. Download the Excel file to sort and filter the protection results.

Exported fields

The following tables list the fields that are exported using the **Export All** functionality. The fields are exported in CSV format.

Table 14. Exported fields

Resource	Exported fields
Jobs	Asset Name, Host, Component Type, Component Description, Schedule Frequency, Job ID, Status, Description, Job Type, Sub Type, Asset Type, Assets, Start Time, End Time, Duration, Next Scheduled, Policy Name, Data Transferred, Storage System, Asset Size, Data Compressed, Average Throughput, Total Compression Factor, Reduction Percentage
Application Agents	Host Name, IP, Registration Status, OS, Agent Type, Current Version, Update Status, Port, Application Version, Created Date, Registered Date, Throttling Status, CPU Throttling
Alerts	Message ID, Details, Recommended Action, Severity, Date, Summary, Category, Status, Component Type, Component Description
Messages Catalog	Message ID, Message, Details, Recommended Action, Severity, Category
Protection Policies	Name, Category, Asset Type, Asset Count, Protected Asset Size, Last Run Status, Violations, State
Resource Groups	Name, Description, Created At, Number of Resources
SLA Compliance	Name, Compliance Type, Policies At Risk, Objectives out of Compliance, Impacted Assets
System Health Issues	Deduction, Issue, Category, Component, Remediation, Date
Users	User/Group Name, Type, First Name, Last Name, Email Address, Roles and Resources, Added Date

The following fields are common to each asset type:

ID, Status, Asset Type, Sub Type, Protection Policy ID, Protection Policy, Protection, Size, Protection Capacity Size, Protection Capacity Time, Last Copy, Network, Protection Rule Name, Resource Group Name

The following table lists the fields that are unique to each asset type.

Table 15. Exported fields for asset types

Resource (asset type)	Exported fields
VMware Virtual Machines	Name, Tags, Operating System, Apps, Disk Excluded, vCenter, Protection Mechanism, ESX Host Name, VM BIOS Uuid, Resource Pool, VM Folder, Data Center
Kubernetes	Namespace, Labels, Age, Cluster, PVCs Excluded, Storage Class Name, Volume Mode, PVC Namespace
Microsoft SQL Server	Name, Protection Engine Flow, Host Type, Host/Cluster/Group Name, Application Server ID, Application Server Name
Oracle	Name, Host/Cluster/Group Name, Host Type, OS Type, Application Server Name, Application Server ID, SID, Data Guard Name, Data Guard Role, Protocol, Backup Technology

Table 15. Exported fields for asset types (continued)

Resource (asset type)	Exported fields
Microsoft Exchange Server	Name, Host/Cluster/Group Name, Host Type, Application Server Name, Application Server ID
SAP HANA	Name, Host/Cluster/Group Name, Host Type, Application Server Name, Application Server ID
File System	Name, OS Type, File System Type, Host Name, Host Operating System
NAS	Name, Asset Source, Appliance Name, Array Type, Server Name/IP, Protocol, File Stubs, File System Path, File System Name
VMAX storage group	Name, VMAX Serial No, Host

Customer feedback

Use the customer feedback feature in the PowerProtect Data Manager UI to report your satisfaction with PowerProtect Data Manager, provide feedback, and send requests for enhancements. Customer feedback is used to improve the customer experience.

Provide general feedback

Use the following procedure to report your satisfaction with PowerProtect Data Manager and provide feedback.

Steps

1. Log in to the PowerProtect Data Manager UI.

2. From the banner, click .

The customer feedback survey opens in a new window.

NOTE: In environments with limited external connectivity, such as dark sites, an error appears in the web browser and the customer feedback survey is not displayed.

3. (Optional) Complete the fields in the customer feedback survey, and when finished, click **Submit**.

You have the option to rate your satisfaction with PowerProtect Data Manager and make a recommendation for how to improve the customer experience. You also have the option to provide an email address so that can follow up with you regarding your feedback.

NOTE: Customer contact information is not used for marketing purposes.

Security configuration

A separate guide provides some server configuration tasks which are intended specifically for PowerProtect Data Manager security administrators, whose role may be separate from the PowerProtect Data Manager host system administrator.

The *PowerProtect Data Manager Security Configuration Guide* provides detailed instructions for all security-related tasks, including but not limited to:

- Port requirements for and between the following components:
 - PowerProtect Data Manager
 - Configured DD systems
 - VM Direct appliances (embedded and external)
 - Application-agent hosts
 - Web and REST API clients
 - Callhome (SupportAssist)

- ESXi
- vCenter
- Configuring identity providers
- Managing local and external user accounts
- Changing and resetting passwords
- Assigning users and groups to roles and associated privileges
- Managing credentials for local and remote components
- Creating resource groups to define scopes of authority
- Managing security certificates, where applicable

Role-based security

PowerProtect Data Manager provides predefined user roles that control access to areas of the user interface and to protected operations. Some PowerProtect Data Manager functionality is reserved for particular roles and may not be accessible from every user account.

By using the predefined roles, you can limit access to PowerProtect Data Manager and to backup data by applying the principle of least privilege.

The *PowerProtect Data Manager Security Configuration Guide* provides more information about user roles, including the associated privileges and the tasks that each role can perform.

System Maintenance

Topics:

- Deploying and maintaining the health of PowerProtect Data Manager
- Deploying and updating PowerProtect Data Manager
- Licensing PowerProtect Data Manager
- Specifying the PowerProtect Data Manager host
- Memory optimization
- Restricted mode
- System support
- Restarting PowerProtect Data Manager
- System maintenance troubleshooting
- Messages Catalog

Deploying and maintaining the health of PowerProtect Data Manager

In order for PowerProtect Data Manager to function as efficiently as possible, you should deploy and maintain it according to recommended guidelines.

Deploying and updating PowerProtect Data Manager

You can deploy PowerProtect Data Manager, update it to the latest version, and install other important package updates.

Update paths

CAUTION: If recommended guidelines are not followed, the update of PowerProtect Data Manager or one of its components can fail.

When deploying or updating PowerProtect Data Manager, see the *PowerProtect Data Manager Deployment Guide*. It contains detailed instructions and guidelines that must be followed in certain environments and configurations.

Updating from PowerProtect Data Manager versions 19.10 through 19.13 to version 19.14 is supported.

Security advisories

CAUTION: If the latest Dell security advisories (DSAs) are not followed, PowerProtect Data Manager can be exposed to security vulnerabilities.

To review the latest DSAs, search for PowerProtect Data Manager at the Dell Technologies Security Advisories and Notices website.

Licensing PowerProtect Data Manager

PowerProtect Data Manager can be licensed in several different ways. This section describes the different types of available licenses and how to install a license.

For more information about licensing, see the *PowerProtect Data Manager Deployment Guide*.

License types

There are several different types of licenses, and they can provide licensing for different periods of time.

The available license types are described in the following table.

Table 16. License types

License type	Description
Trial	The license used by default when PowerProtect Data Manager is deployed. It enables full use of the product without adding a license key for up to 90 days. When the trial period ends, PowerProtect Data Manager continues to operate with full functionality so that you can add a permanent license. NOTE: A trial license does not allow the use of SupportAssist.
Front-end protected capacity by terabyte (FETB)	The primary model of licensing, which is based on the capacity that you want to protect. For example, you can purchase a 100-TB license, which enables you to protect up to 100 TB of data.
Socket-based	Licensed per CPU socket on virtual machine hosts that are being backed up or replicated.

Perpetual and term-based (subscription) licenses

Licensed software is offered with perpetual or term-based licenses. Your quote identifies whether your license rights are perpetual or term-based.

A perpetual license enables you to use the software while you are in compliance with the terms of the license agreement.

A term-based license enables you to use the software for a specified time, while you are in compliance with the terms of the license agreement. At the end of the license term, you must stop using the software, extend the license term, or purchase a new license.

Add a license

You can add a license file to PowerProtect Data Manager and view license details, such as capacity usage and software ID number.

Prerequisites


To obtain the XML license file from the license management website, you must have the License Authorization Code (LAC), which is emailed from . If you have not received the LAC, contact your Customer Support representative.

About this task

To review existing license information, go to **Settings > License**.

To add a license, perform the following steps:

Steps

1. From the PowerProtect Data Manager user interface, click , and then select **License**.
2. On the **License** window, perform one of the following actions:

- Copy and paste the text from the license file into the text box.
- Click **Upload File**, browse to the location of the license file and select the file, and then click **Open**.
The license file content appears in the **License** window.

3. Click **Apply**.

Results

A message appears in the **License** window to confirm that the license is successfully added.

Specifying the PowerProtect Data Manager host

When you specify a vCenter server as the PowerProtect Data Manager host, it allows the vCenter server to perform operations unique to PowerProtect Data Manager.

The PowerProtect Data Manager host performs several operations, including the following:

- Virtual-machine configuration and other system activities.
- Taking a PowerProtect Data Manager snapshot, if required during a software update.
- Allowing memory that is assigned to PowerProtect Data Manager to be automatically increased as necessary when performing a software update.
- Enabling Cloud Disaster Recovery (Cloud DR) in order to increase required PowerProtect Data Manager CPU and memory. A vCenter host is a prerequisite for Cloud DR, as specified in the **Cloud Disaster Recovery** tab of the **Infrastructure > Asset Sources** window in the PowerProtect Data Manager user interface.

Specify a vCenter server as the PowerProtect Data Manager host

You select a vCenter server to be used as the PowerProtect Data Manager host from those already added or discovered.

About this task

Perform the following operations:

Steps

1. From the PowerProtect Data Manager user interface, click  and then select **Hosting vCenter**.
The **Hosting vCenter** window appears.
2. Choose from one of the following options:
 - **Enter FQDN/IP**—Select this option to manually enter the fully qualified domain name or IP of the vCenter server, the port number, and to select the vCenter **Host Credentials**. The **Host Credentials** list is populated with vCenter servers that have already been added and discovered in PowerProtect Data Manager. If the host vCenter credentials do not appear in the list, select **Add Credentials** to enter this information.
 - **Select FQDN/IP from asset sources**—Select this option to obtain the host vCenter server information automatically from a vCenter asset source that has already been added and discovered in PowerProtect Data Manager.
3. Click **Save**.

Results

If the host vCenter server is added as an asset source in PowerProtect Data Manager,  is displayed next to this vCenter server in the **Infrastructure > Asset Sources** window.

Minimum privileges required for the vCenter server PowerProtect Data Manager host

The user account associated with the vCenter server that is specified as the PowerProtect Data Manager host must have the following minimum privileges. These privileges are required for functions related to software installation and updates, virtual machine snapshots and rollback, and configuring virtual machine memory.

Setting	vCenter 6.0 and later required privileges	PowerCLI equivalent required privileges
Global	<ul style="list-style-type: none"> Manage custom attributes Set custom attributes 	<ul style="list-style-type: none"> Global.ManageCustomFields Global.SetCustomField
Network	<ul style="list-style-type: none"> Assign network 	<ul style="list-style-type: none"> Network.AssignNetwork
Permissions	<ul style="list-style-type: none"> Modify permission 	<ul style="list-style-type: none"> Authorization.ModifyPermissions
Sessions	<ul style="list-style-type: none"> Impersonate user Message Validate session View and stop sessions 	<ul style="list-style-type: none"> Sessions.ImpersonateUser Sessions.Message Sessions.ValidateSession Sessions.ViewandStopSessions
Virtual Machine	<ul style="list-style-type: none"> Change Configuration > Add or remove device Change Configuration > Change CPU count Change Configuration > Change Memory Change Configuration > Change Settings 	<ul style="list-style-type: none"> VirtualMachine.Config.AddorRemoveDevice VirtualMachine.Config.CpuCount VirtualMachine.Config.Memory VirtualMachine.Config.Settings
Virtual Machine	<ul style="list-style-type: none"> Snapshot Management > Create snapshot Snapshot Management > Revert to snapshot Snapshot Management > Remove snapshot Snapshot Management > Rename snapshot 	<ul style="list-style-type: none"> VirtualMachine.State.CreateSnapshot VirtualMachine.State.RevertToSnapshot VirtualMachine.State.RemoveSnapshot VirtualMachine.State.RenameSnapshot

NOTE: A complete list of the privileges required for a dedicated vCenter user account is provided in the *PowerProtect Data Manager Virtual Machine User Guide*.

Memory optimization

You can use adjust the amount of memory that is assigned to the PowerProtect Data Manager virtual machine in order to optimize server performance.

The following table indicates the default amount of memory assigned to the PowerProtect Data Manager virtual machine in a standard environment. The default values are the minimum recommended values.

Table 17. PowerProtect Data Manager memory requirements

Deployment type	Memory	Swap space	Cores
Default	24 GB	8 GB	10
With the Cloud Disaster Recovery (Cloud DR) Add-On	28 GB	8 GB	14

The recommended number of cores is 14. Also consider the following:

- Depending on the environment, increasing the amount of memory can increase performance.
- If low-memory alerts are seen, increase the amount of memory.
- Do not increase the amount of memory beyond 32 GB of RAM. PowerProtect Data Manager is not designed to support more than 32 GB of RAM.
- If you are deploying PowerProtect Data Manager to a virtual machine in a cloud Marketplace environment, it is automatically assigned 52 GB of RAM. This amount of memory should not be changed after it is deployed.
- Most of the services from PowerProtect Data Manager are memory intensive. When the available physical memory drops to a certain threshold value, these services start leveraging swap memory. If swap memory resides on a slow disk, then there

can be significant impact on the Java Garbage Collection activity from each of these services when memory that has not been recently used needs to be swapped into physical memory.

- It is highly recommended to configure swap memory on a solid-state drive (SSD). During deployment of the PowerProtect Data Manager server, use the SSD data store to avoid the high latency disk impact from swap and metadata operations.

NOTE: For help with optimizing memory, contact your Customer Support representative.

Memory and updating from an earlier version of PowerProtect Data Manager

Features in the current version of PowerProtect Data Manager might require more memory than required in previous versions. When updating from an earlier version of PowerProtect Data Manager, ensure that you increase the amount of assigned memory as necessary.

Adjust the virtual machine memory

Adjust the amount of memory assigned to the PowerProtect Data Manager virtual machine to support changes in the protection environment.

Steps

1. Log in to the **vSphere Web Client**.
2. Right-click the appliance and select **Edit Settings**.
The **Edit Settings** window appears with the **Virtual Hardware** button selected.
3. In the **Memory** field, specify the new memory value.
Ensure that the value you specify does not exceed 32 GB of memory and that it is a multiple of 4 GB.
4. Click **OK**.

Restricted mode

You can enable restricted mode to prevent scheduled writes to storage. You might want to enable restricted mode to limit access to storage during a storage upgrade.

Enabling restricted mode during a storage upgrade provides the following benefits:


- Storage writes can be eliminated in a controlled manner. Once writes have stopped, storage can be upgraded.
- Storage writes can be tested after storage has been upgraded. Once testing is complete, storage can be returned to full production.

Restricted mode prevents the following scheduled operations:

- Backups and replication
- Backup-copy deletion
- Server disaster-recovery backups


Restricted mode allows the following operations:

- Any jobs in progress or queued to run
- Manual backups and restores
- Discovery jobs

To enable restricted mode from the PowerProtect Data Manager user interface, click  select **Support > Restricted Mode**, and then click **Enable Restricted Mode**.

System support

You can use the PowerProtect Data Manager user interface to manage and modify support settings that are typically configured during deployment. Typically configured support settings include the mail server setup and Secure Remote Services registration.

To access the **Support** window, click , and then select **Support**.

Configuring SupportAssist for PowerProtect Data Manager

SupportAssist is a support tool that communicates with PowerProtect Data Manager to monitor your environment, automatically detect current and potential issues, and collect and store diagnostic data. SupportAssist securely sends the data that is required for troubleshooting an issue to Customer Support for diagnostic purposes and customer support.

SupportAssist is at heart of the connectivity platform as a unified communication point between PowerProtect Data Manager and Customer Support.

SupportAssist provides the following features and benefits:

- Proactive monitoring and issue prevention
- Facilitates update package downloads
- Automatic health checks
- Communicates telemetry data
- Real-time troubleshooting
- Customer support

Configure SupportAssist to receive automated support capabilities for your PowerProtect Data Manager system.

SupportAssist cannot be configured when PowerProtect Data Manager uses a trial license.

Generate SupportAssist access key and PIN

An access key and PIN are required to configure a secure connection between PowerProtect Data Manager and SupportAssist. You only need to apply the access key and PIN once.


About this task

Use the following procedure to generate your SupportAssist access key and PIN:

Steps

1. Go to the Customer Support website and log in to your account.
2. In the search box, type PowerProtect Data Manager and click **Search**.
3. Click **Generate Access Key** in the **Quick links** pane.
4. Enter the product ID (serial number) in the search box.
5. In the **Create PIN** field, enter a 4-digit PIN.
Record the PIN for later use.
6. Click **Generate Access Key**.

The access key is sent to the email address for your account.

 **NOTE:** It might take up to 5 minutes to receive the access key in your email.

Connect to support services through SupportAssist

Establish a connection through SupportAssist to ensure access to Customer Support. SupportAssist enables PowerProtect Data Manager to connect to support services directly or through a gateway server.

Prerequisites

- Apply a valid PowerProtect Data Manager license.

- If you are connecting through the gateway server, the SCG gateway version must be 5.10 or later.
- Apply a valid access key and PIN.
- HTTPS port 443 of `esrs3-core.emc.com` and `esrs3-core-dr.emc.com` is not blocked by the network firewall.

Steps

1. From the PowerProtect Data Manager UI, click , select **Support**, and then click **SupportAssist**. The **Support** window opens to the **SupportAssist** page.
2. On the **Connection** tab, click **Connect Now**.
3. Select one of the following options:
 - **Connect Directly**
Select this option to connect PowerProtect Data Manager directly.
 - **Connect via Gateway**
Select this option to connect PowerProtect Data Manager through a gateway server, and then enter the gateway server IP address and port number.
4. Enter the SupportAssist Access Key and PIN.
5. Click **Enable Connect**.

Results

PowerProtect Data Manager is connected to support services.

Update or configure contact data

Provide contact information for the person that Customer Support will contact with diagnostic reports. You can add or update contact data for SupportAssist at any time.

Steps

1. From the PowerProtect Data Manager UI, click , select **Support**, and then click **SupportAssist**. The **Support** window opens to the **SupportAssist** page.
2. Select the **Contacts** tab.
3. To add a primary contact, complete the following steps:
 - a. Enter the following information:
 - **First Name**
 - **Last Name**
 - **Email**
 - **Phone**
 - b. Select the **Preferred Language** from the list.
 - c. Click **Save**.
4. To add a secondary contact, click **+ Add Secondary Contact** and enter the required information.

Change SupportAssist connection settings

Use the following procedure to change SupportAssist connection settings.

Steps


1. From the PowerProtect Data Manager UI, click , select **Support**, and then click **SupportAssist**. The **Support** window opens to the **SupportAssist** page.
2. Select one of the following connection options:
 - **Connect Directly**
 - **Connect via Gateway**

To add a new gateway connection, complete the following steps:

- a. Enter the gateway IP address and port number.
- b. Click **Test**.

Wait until the connection test is complete. If the connection is successful, a green check mark is displayed next to the gateway IP address and port number.

3. Enter the SupportAssist Access Key and PIN.


 **NOTE:** If you are not connecting with a new access key, skip this step.

4. Click **Reconnect**.

Enable or disable SupportAssist

Enable the SupportAssist feature to automatically detect issues and collect diagnostic and usage data. You can also disable SupportAssist at any time.

Steps

1. From the PowerProtect Data Manager UI, click , select **Support**, and then click **SupportAssist**. The **Support** window opens to the **SupportAssist** page.
2. To enable SupportAssist, move the **Connect to SupportAssist** slider to the right. To disable SupportAssist, move the **Connect to SupportAssist** slider to the left.
The operation might take up to 5 minutes to complete.

Troubleshooting SupportAssist

Review the following information that is related to troubleshooting SupportAssist.

Failed to establish a SupportAssist connection

If you are connecting to SupportAssist with an access key and PIN that is already in use, the connection fails with error:

`Connection is failed: Get universalkey error: Access Key and Pin used`

If this issue occurs, obtain a new access key and PIN from Customer Support. Generate SupportAssist access key and PIN provides instructions.

The following error might display if the SWID is not added to the PowerProtect Data Manager back-end: `Connection is failed: Get universalkey error: Invalid Access Key and Pin`

If this issue occurs, contact Customer Support and ask them to check whether the SWID has been added to the PowerProtect Data Manager back-end.

Connection status changes to "Not Connected"

If the connection status changes to "Not Connected":

1. Ensure that all prerequisites are met in Connect to support services through SupportAssist .
2. If the issue persists, contact Customer Support.

Telemetry Collector

Telemetry Collector gathers information related to this system, including configuration, usage characteristics, performance, and deployment location information. Telemetry Collector manages remote access and the exchange of system data with Dell Inc. or its subsidiaries. The information that is gathered by Telemetry Collector is confidential and this data cannot be shared.

When you enable SupportAssist, you also enable Telemetry Collector, which allows Customer Support engineers to collect data that is related to troubleshooting device and PowerProtect Data Manager software issues. Telemetry Collector does not collect any personal information.

Telemetry Collector populates three reports—a telemetry report, an alert summary report, and a CloudIQ report. Telemetry Collector collects details about the following objects:


- Alerts
- Assets
- Asset sources
- Audit logs
- Cloud Data Recovery
- Cloud Disaster Recovery metrics
- Compliance details
- Compliance in the last 24 hours
- Data targets
- DD inventory
- Host information
- Integrated storage
- Licensing
- PowerProtect Data Manager operational inventory
- Protection details
- Protection policies
- Quick-recovery synchronization information
- Service-level agreements
- Storage systems
- Time spent on generating reports
- Traffic metrics
- Update summaries
- Usage

CloudIQ reporting

When you enable AutoSupport and choose SupportAssist, you also enable reporting. CloudIQ is a no-cost SaaS/cloud-based management application that proactively monitors and measures the overall health of systems through intelligent, comprehensive, and predictive analytics. The data reported to CloudIQ includes configuration data, historical metrics and health score data.

Ensure that the following requirements are met:

- Add a valid license in **System Settings > License**.
- Set up SupportAssist in **System Settings > Support > SupportAssist**.
- Enable AutoSupport and select **SupportAssist**.

When AutoSupport is enabled, CloudIQ reports are sent automatically. To log in to CloudIQ, click , and then click CloudIQ. You can also go to <https://cloudiq.dell.com>. For more information on CloudIQ, refer to the CloudIQ Online Support site.

Set up the email server

The **Email Setup** page of the PowerProtect Data Manager **Support** window enables you to configure SMTP email server settings that control sending and receiving email related to resetting local user passwords and customizing alert notifications.

Steps

1. From the PowerProtect Data Manager user interface, click , select **Support**, and then click **Email Setup**.
2. Populate the following fields:
 - a. **Mail Server**
The SMTP mail server.
 - b. **Email from:**
The email address at which you would like to receive PowerProtect Data Manager AutoSupport email.
 - c. [Optional] **Recipient for Test Email:**

The email address to which you would like to send PowerProtect Data Manager test email.

d. [Optional] **Port:**

The default port is 25. PowerProtect Data Manager supports using non-default ports.

If the email setup is deleted, you must manually choose any non-default port that is not in use anywhere else.

e. **User Name:**

The user name associated with the PowerProtect Data Manager SMTP email server. This field is optional.

f. **Password:**

The password associated with the PowerProtect Data Manager SMTP email server. This field is optional.

3. Click **Send Test Email**.

PowerProtect Data Manager sends a test email.

4. Click **Save**.

Add AutoSupport

When AutoSupport is enabled, automated support information, telemetry reports, alert summaries, and CloudIQ reports are sent.

About this task

If SupportAssist and SMTP are both configured, this information is sent using the option that you choose in the **System Settings > Support > AutoSupport** window.

Steps

1. From the PowerProtect Data Manager UI, click , select **Support**, and then click **AutoSupport**.

The **AutoSupport** window appears.

2. Change the **Enable AutoSupport** option to **Disabled** or **Enabled**, and click **Save**.

When you enable AutoSupport, select whether to receive the AutoSupport communications through SupportAssist or email server.

When you enable AutoSupport, the **Telemetry Software Terms** page displays. Review and scroll down to the bottom of the page to accept the terms, and then click **Save** to save your changes.

When you disable AutoSupport, PowerProtect Data Manager stops sending error and telemetry data to SupportAssist or the SMTP server. PowerProtect Data Manager continues to send information for updates and other information.

 **NOTE:** To disable SupportAssist, clear the SupportAssist option in the AutoSupport window.

Enabling automatic update package checks and downloads


If SupportAssist is enabled, you can configure PowerProtect Data Manager to automatically check for update packages, and either alert you or automatically download them.

For more information about these options, see the *PowerProtect Data Manager Deployment Guide*.

Add a log bundle

Use the following procedure to add a log bundle.

About this task

 **NOTE:** You can add a maximum of 10 log bundles.

Steps

1. From the PowerProtect Data Manager user interface, click  and then click **Logs**.

2. Click **Add** to add a log bundle.
The **Add Log Bundle** window appears.
3. Select the systems for the log bundle (**Data Manager**, **VM Direct Engines**, or, if Cloud DR is deployed, **CDRS**), set the log bundle duration, and click **Save**.
The **Jobs** window displays the progress of the log bundle creation. Also, a green banner in the UI indicates that the log bundle has successfully been created. If you want to dismiss the banner, click **X**.
4. To delete the log bundle, select the box to the left of log bundle and click **Delete**.
The **Log Capacity** indicates how much space (in GB) remains on the disk for logs and the percentage of the disk in use for log storage.
5. To download the log bundle, click the bundle name in the **Bundle Name** column.

Audit logging and monitoring system activity

The Linux audit daemon (`auditd`) tracks and logs security-relevant events on the PowerProtect Data Manager system.

Users with the Administrator role can use `auditd` to monitor the following events:

- File access
- System calls
- Login and logout activity of users

Audit logging enables you to discover access violations, changed or deleted files, failed authentication, and so on.

Viewing audit events in the UI

With the Administrator, Backup Administrator, Restore Administrator, and User roles, you can view audit events to monitor system activity.

About this task

The following actions generate an audit event:

- User login and logout
- Creating, deleting, or updating a user
- Assigning or unassigning a role to a user

To view audit events in the UI, perform the following steps.

Steps


1. Log in to the PowerProtect Data Manager UI with an account that has one of the indicated roles.
2. Go to **Administration > Audit Logs**.

View and manage alerts

Alerts enable you to track the performance of data protection operations in PowerProtect Data Manager so that you can determine whether there is compliance to service level objectives. With the Administrator, Backup Administrator, Restore Administrator, or User role, you can access the alerts from the **Alerts** window. However, only some of these roles can manage alerts.

Steps



1. From the PowerProtect Data Manager UI left navigation pane, select **Alerts**.


You can also click  on the top banner, and then click the links to view unacknowledged alerts of all statuses (critical, warning, and informational), or only the unacknowledged critical alerts.

 **NOTE:** Clicking the **New** tag displays only the unacknowledged alerts that have been generated within the last 24 hours.

The number that appears next to  is the total number of unacknowledged critical alerts over the last 24 hours.

The **Alerts** window displays.


2. Select the **System** tab. A table with an entry for each applicable alert displays.
By default, only unacknowledged critical alerts from the last 24 hours display, unless you selected to view all unacknowledged alerts from the links under .
If filter tags have already been applied, the window displays these filter tags. Click **X** next to any of these filter tags to clear a filter, and the table view updates with the applicable selections. You can sort the alerts in the table by Severity (Critical, Warning, Informational), Date, Category, or Status (Acknowledged or Unacknowledged).
3. Select a time from the last 24 hours, the last 3 days, the last 7 days, the last 30 days, and a specific date for the alerts you want to view, or provide a custom time range. You can also select **All Alerts** from this list to display information for all alerts that match the filter tags.
4. Optionally, clear the **Show only unacknowledged alerts** checkbox if you want to view both acknowledged and unacknowledged alerts. If you clear this checkbox, the **Unacknowledged** filter tag is also cleared.
5. To view more details about a specific entry, click  next to the entry in the table.
6. For the following steps, log in to the PowerProtect Data Manager UI with an account that has the Administrator, Backup Administrator, or Restore Administrator role.
7. To acknowledge one or more alerts, select the alerts and then click **Acknowledge**.
8. To add or edit a note for the alert, click **Add/Edit Note**, and when finished, click **Save**.
9. To export a report of alert information to a .csv file which you can download for Excel, click **Export All**.

 **NOTE:** If you apply any filters in the table, exported alerts include only those alerts that satisfy the filter conditions.

Export audit logs

With the Administrator or Security Administrator role, you can export audit log records to a .csv file of audit data that you can download and open in Excel. Only the Administrator role can change the retention period.

Steps

1. Go to **Administration > Audit Logs**.
The list of audit logs appears, which displays the following information:
 - Changed at
 - Audit Type
 - Description
 - Changed By
 - Object Changed
 - Previous Values
 - New Values
2. To set the retention period (in days) for the audit log, select **Set Boundaries** and update the retention period.
Only the Administrator role can perform this step.
3. To add a note for the audit log, click , enter a note in the **Notes** field, and click **Save**.
4. Click **Export All**.

Monitor system services and system health

The status of system services can be monitored from the **System Services Status** pane, and system health information can be monitored from the **Health** pane.

Monitor system services

You can monitor the status of each system service from the **System Services Status** pane.

To view the status of system services, click , select **Support**, and then click **System Services Status**.

The following table provides a summary of the status of each system service and component:

Table 18. System service and component status

Status	Description
Running	This state appears when the associated service or component is running with full functionality. When all services are in running state, the state of the appliance is operational.
Initializing	This state appears when the service is starting. When the service successfully starts, the state changes to Running.
Maintenance	This state appears when the associated service is in maintenance. In the maintenance state, components have limited functionality. Infrastructure services do not go into maintenance state. When other services or components are in maintenance, the appliance state is also maintenance.
Quiesce	This state appears when the service or service associated with the component is stopping.
Shut down	This state appears when the service has stopped.
No response	This state appears when the service that is associated with the component is running, but the service is not responding.

Monitor system health

You can monitor system health information from the **Health** window of the PowerProtect Data Manager UI.

To view a summary of any issues affecting the health of PowerProtect Data Manager, select **Health** from the navigation pane or **View All** from the Dashboard health widget.

PowerProtect Data Manager automatically performs a health check every two minutes. If an issue is detected, it is assigned a category and a deduction value based on its severity. All issues are displayed on the **Health** window. Resolved issues are automatically removed the next time a health check is performed.

Health details and status are provided for the following categories:

- **Components** identifies the state of hardware and software services, such as Running or Failed.
- **Configuration** identifies whether any aspects of the PowerProtect Data Manager configuration are incomplete, such as System Support configuration.
- **Capacity** identifies the provisioned and currently allocated size of the associated storage system.
- **Performance** identifies key performance indicators, such as memory use.
- **Data Protection** identifies key protection indicators, such as service-level agreements not being met and disaster-recovery backup copies not being present.

Each category starts with a score of 100. If there is an outstanding health check issue in one of these categories, its score is reduced by the deduction value assigned to the issue. If there is more than one outstanding issue in the category, its score is only reduced by the deduction value of the most severe issue.

Click  next to an entry to see the details of the issue.

In the **Health** window, you can export health data by using the **Export All** functionality.

The overall health score of the system is represented by the most severe issue and the category with the lowest score.

Table 19. Overall health score

Health score	Indicates
95–100	System is in good health.
71–94	System is in fair health.
0–70	System is in poor health.

Table 20. Health check descriptions

Category	Health Check	Maximum Deduction	Description
Configuration	Asset source configuration	-30	Deduction occurs when no asset sources have been added and enabled in PowerProtect Data Manager. When at least

Table 20. Health check descriptions (continued)

Category	Health Check	Maximum Deduction	Description
			one asset source is added, the health score returns to normal.
	Storage configuration	-30	Deduction occurs when there are no storage targets configured in the system. When at least one storage target is set, the health score returns to normal. ① NOTE: Not applicable to the PowerProtect Data Manager appliance.
	Support configuration	-10	Deduction occurs when there are no support options configured in the system. The support options include: <ul style="list-style-type: none"> • Email setup • Support assist • Auto support When a support option is configured, the health score returns to normal. If a support option is configured but initialization is still in progress, the health score reduction is set to -5 until the initialization is complete.
	Policies defined for all assets	-2	Deduction occurs if any of the assets from the asset sources enabled in PowerProtect Data Manager are not protected (for example, Protected/Exclude). The deduction is -2 when unprotected assets total is greater than 0. Once all assets have been moved to a protected state, the health score returns to normal.
	System disaster recovery (DR) backup schedule	-10	Deduction occurs when there is no scheduled system DR backup. Once the system DR schedules have been set, the health score returns to normal.
	License	-30	Deduction occurs when the license status is not valid or close to its expiration date: <ul style="list-style-type: none"> • When the license is invalid or has expired: -30 • When the license expires in less than 7 days: -20 The health returns to normal upon application of a valid PowerProtect Data Manager license.
	Operating system account health check	-60	Deduction occurs if any of the operating system account passwords are about to expire or already expired: <ul style="list-style-type: none"> • Before operating system account password expiry: -15 • Upon password expiry: -60

Table 20. Health check descriptions (continued)

Category	Health Check	Maximum Deduction	Description
			Once the operating system account expiry error is fixed, the health score returns to normal.
	Search cluster configuration	Search cluster is disabled: -5 Search node parent vCenter Servers are removed: -5	Deduction occurs when the Search cluster is disabled or the parent vCenter Servers are removed. NOTE: Not applicable to the PowerProtect Data Manager appliance. The health score returns to normal once the Search cluster is properly configured.
	Reporting cluster configuration	When reporting node parent vCenter Servers are removed: -5	Deduction occurs when the Reporting node parent vCenter Servers are removed. NOTE: Not applicable to the PowerProtect Data Manager appliance. Once the reporting cluster error is fixed, the health score returns to normal.
	ES configuration	-5	Deduction occurs when undefined ES settings have been added. Once the error is fixed on the ES side, the health score returns to normal.
Components	PowerProtect Data Manager core infrastructure services status	Business services: 30 Core services: 30 Infrastructure services: 60 Management services: 40 Protection services: 20	Deduction occurs when one or more of the PowerProtect Data Manager services is not running or is disabled. The health score returns to normal when all services are up and running.
	Protection engines status	-10	Deduction occurs when the protection engine requires attention. The health score/status returns to normal when the protection engine status is in operational state.
	Reporting	-10	Deduction occurs when one or more of the report nodes cannot be detected. Once the health check error is fixed, the health score returns to normal.
	Search cluster	-25	Deduction occurs when one or more Search clusters or nodes are disabled or cannot be detected. The health score/status returns to normal once all the Search cluster issues are resolved.
	Cloud Disaster Recovery	-25	Deduction occurs when the Cloud DR Server in PowerProtect Data Manager cannot be detected or the password is invalid.

Table 20. Health check descriptions (continued)

Category	Health Check	Maximum Deduction	Description
			The health status/score returns to normal once the DD Cloud Disaster Recovery server issues have been resolved.
	Heap dump	-2	Deduction occurs when java heap dump files are detected in the java service log folder. The health score returns to normal when there are no java heap dump files detected.
	DNS	-60	Deduction occurs when all the DNS servers are unreachable. The health score returns to normal when at least one of the DNS servers can be reached.
	NTP	-10	Deduction occurs when all the NTP servers are unavailable. The health score returns to normal when at least one of the configured NTP servers can be reached.
	ES Shards Health Check	-50 (replica shards unassigned) -70 (primary shards unassigned)	Deduction occurs when the Replica or Primary shards are unassigned. Once the ES Shards errors are fixed, the health score returns to normal.
Data Protection	Service Level Agreement (SLA) compliance	-50	Deduction occurs when SLA compliance is defined but has not been met, for example, asset compliance ratio is defined as: $\text{Out Of Compliance Asset Count} / \text{In Compliance Asset Count} + \text{Out Of Compliance Asset Count}$ <ul style="list-style-type: none"> Low ratio: Compliance ratio $\leq 1/3$ High ratio: $1/3 < \text{Compliance ratio} \leq 2/3$ Critical ratio: Compliance ratio $> 2/3$ <p>When more than 2/3 of protection policies are out of compliance with the defined SLAs, the score deduction is -50. The health score returns to normal when the SLA compliance has been met, for example, <code>complianceRatio= 0</code>.</p>
	System DR backup copy present	-40	Deduction occurs when the System DR backup copy is not present. When the DR backup copy exists, the health score returns to normal.
	Discovery status	-20 (for PowerProtect Data Manager) -5 (for Cloud Snapshot Manager)	Deduction occurs when the PowerProtect Data Manager or Cloud Snapshot Manager (CSM) discovery job completes with an error. The health score returns to normal once the Discovery jobs errors are fixed.
Capacity	PowerProtect Data Manager disk space	-60	Deduction occurs when there is heavy disk partition space use. When disk space usage is 75-90%, the score deduction is -15. When the disk space usage exceeds 90%, the score deduction is -60.

Table 20. Health check descriptions (continued)

Category	Health Check	Maximum Deduction	Description
			The health score returns to normal when disk space usage falls below the 75% threshold.
Performance	Memory usage	-40	Deduction occurs when there is heavy operating system memory usage. When memory usage is 80-9%, the score deduction is -15. When the memory usage exceeds 95%, the score deduction is -40. The health score returns to normal when disk space usage falls below the 80% threshold.

The following health checks provide grace periods, allowing you a period of time after deployment to configure your system without a significant reduction in the overall health score. An informational alert notification appears up to 24 hours before the score deduction occurs.

Table 21. Deductions with grace period

Health check component	Deduction by grace period
Asset source configuration	<ul style="list-style-type: none"> Not configured for up to 48 hours: -5 Not configured for more than 48 hours but less than one week: -20 Not configured after more than one week: -30
Storage configuration	<ul style="list-style-type: none"> Not configured for up to 24 hours: -5 Not configured after 24 hours: -30
System Support configuration	<ul style="list-style-type: none"> Not configured for up to one week: -5 Not configured after more than one week: -10
System Disaster Recovery (DR) backup schedule	<ul style="list-style-type: none"> Not configured for up to 48 hours: -5 Not configured for more than 48 hours: -10

Access the open source software package information

All open source software (OSS) package information used by PowerProtect Data Manager is stored in a common directory.

To access this information, SSH login to PowerProtect Data Manager and retrieve the OSS reports from the `/usr/local/brs/puppet/licenses` directory.

Security certificates

A default deployment of PowerProtect Data Manager creates self-signed security certificates that secure communication with other components. As you configure the server and add assets, PowerProtect Data Manager stores additional certificates for each component.

The Administrator and Security Administrator roles can review the **Administration > Certificates** page in the UI. This page contains three tabs that list the installed security certificates. Each tab provides information about certificate uses, expiry dates, issuers, and so forth.

The certificates on the **Internal** tab secure access to components that are part of the PowerProtect Data Manager server, such as the UI and REST API. The certificates on the **Application Agents** tab secure access to the agents, which are under the control of PowerProtect Data Manager but exist outside the server. The certificates on the **External Servers** tab secure access to components or systems that are beyond the control of the server, but where you have approved the communication.

The *PowerProtect Data Manager Security Configuration Guide* contains more information about cryptography and security certificates. This guide provides instructions for how to manage the installed certificates, including important prerequisites, operational considerations, associated tasks, and troubleshooting. For example, you can replace the default self-signed security certificates for PowerProtect Data Manager with certificates from an approved certificate authority. This guide also contains instructions for establishing certificate-based trust with external components and systems.

Restarting PowerProtect Data Manager

When a PowerProtect Data Manager restart is required, it is recommended that you avoid directly powering off the virtual machine unless it is necessary.

To ensure that PowerProtect Data Manager is able to properly restart, use the `reboot` or `shutdown` command. For example, on Linux, run the command `shutdown -r` or `shutdown -h now`.

System maintenance troubleshooting


Services do not start after restarting PowerProtect Data Manager

If the operating system root password expires and you do not change the password before you restart PowerProtect Data Manager, some scripts fail to obtain root privileges. In this situation, the PowerProtect Data Manager services cannot start.

Follow the guidance in the *PowerProtect Data Manager Security Configuration Guide* for operating system expired password behavior to change the root password. Then, restart PowerProtect Data Manager again.

Messages Catalog

The Messages Catalog in the PowerProtect Data Manager UI provides a list of all informational, warning, and critical messages that PowerProtect Data Manager generates. The message details and recommended action can be used to troubleshoot issues, and the message ID is provided for reference when contacting Dell Customer Support.

From the PowerProtect Data Manager UI, click  and select **Messages Catalog** to view the entire catalog. You can sort the information by each column, or filter the list to view messages that match a specific criteria:

- The **Message ID**, **Message**, **Details** and **Recommended Action** columns allow you to search for text and display only results matching the search text.
- The **Category** and **Severity** columns allow you to select from one or more available options to display only messages that match your selections.

To export all messages or a filtered list of messages as a `.csv` file, click **Export**.

Managing Storage

Topics:

- Protection storage
- Storage units
- Differences in storage system and storage unit space reporting
- Monitoring storage capacity thresholds

Protection storage

Protection storage is the set of configured storage systems where PowerProtect Data Manager stores backup copies, replicated copies, and other important information. Protection storage can include any of the following:

- A DD system, including High Availability PowerProtect DD mode
- An instance of PowerProtect DD Management Center (DDMC) that manages multiple DD systems
- A DDMC Smart Scale system pool

NOTE: Data Domain (DD) is now PowerProtect DD. References to Data Domain or Data Domain systems in this documentation, in the user interface, and elsewhere in the product include PowerProtect DD systems and older Data Domain systems.

The most up-to-date software compatibility information for PowerProtect Data Manager is provided by the E-Lab Navigator.

Observe the following information before you configure protection storage:

- Adding and configuring protection storage requires the Administrator role.
- You cannot add protection storage that runs incompatible versions of DDOS.
- You can only add the same protection storage system once, whether you specify the hostname, FQDN, or IP address.
- You cannot add a PowerProtect DD Management Center instance which has no managed DD systems.
- The first time that you add protection storage, PowerProtect Data Manager automatically configures and enables server DR. The first protection storage system is the default target. System recovery for server DR provides more information.
- Adding protection storage by hostname or FQDN provides maximum flexibility for future IP address changes. PowerProtect Data Manager uses DNS to resolve hostnames and FQDNs when you select these entries for the Management network interfaces. Should you later change the DNS mapping, PowerProtect Data Manager resolves the new address and directs Management communication there. Communication with the Data network is by IP address.

Protection storage is further divided into logical groupings that are called storage units, which hold related data and apply more detailed configuration options.

Click  to open the **Details** pane and see more information about an existing protection storage system.

NOTE: Adding a PowerProtect DD Management Center instance is not required for the Storage Direct agent.

PowerProtect DD Management Center automatic discovery

When you add an instance of PowerProtect DD Management Center, PowerProtect Data Manager automatically discovers all the supported DD systems which that PowerProtect DD Management Center instance manages.

PowerProtect Data Manager displays the discovered DD systems on the **Protection Storage** tab of the **Infrastructure > Storage** window after discovery finishes. It may take a few minutes for the discovered systems to appear.

For each DD system, the **Managed By** column in the table indicates the PowerProtect DD Management Center instance that manages the DD system.

If you add a DD system directly to PowerProtect Data Manager, the **Managed By** column displays the name that you provided for the DD system.

High Availability PowerProtect DD support

PowerProtect Data Manager supports DD systems with High Availability (HA) enabled. The Active-Standby configuration provides redundancy in the event of a system failure. HA keeps the active and standby systems synchronized, so that if the active node were to fail, the standby node can take over services and continue where the failing node left off.

When an active High Availability PowerProtect DD system fails over to its standby High Availability PowerProtect DD system, all in progress PowerProtect Data Manager operations including backup, restore, replication, and Cloud Tier continue unaffected.

To add a High Availability PowerProtect DD configuration as a storage target in PowerProtect Data Manager, select **Infrastructure > Storage** in the PowerProtect Data Manager UI. Add protection storage provides more information.

Virtual machine application-aware protection are only be supported with DDOS version 7.0 or later for HA. The most up-to-date software compatibility information for PowerProtect Data Manager is provided by the E-Lab Navigator.

For details on DD systems with HA enabled, see the *DDOS Administration Guide*.

Smart Scale system pools

A system pool is a logical group of DD systems with one interface to flexible storage options. PowerProtect Data Manager can use a system pool as protection storage.

The *DDOS Administration Guide* and *PowerProtect DD Management Center: Installation and Administration Guide* provide more information about Smart Scale, system pools, and the available features. The DDMC instance must be Smart Scale-enabled to use system pools.

After you add the DDMC instance, PowerProtect Data Manager automatically discovers any available system pools. The **Model** column on the **Protection Storage** tab indicates that the protection storage system is a system pool.

PowerProtect Data Manager groups system pools under a separate heading in the list for protection storage selection when working with protection policies.

NOTE:

Adding a DDMC instance with system pools also discovers the individual systems within the system pool. PowerProtect Data Manager includes these systems in lists of available storage targets, such as for protection policy creation. As with a non-Smart Scale DDMC instance, the **Infrastructure > Storage** page groups and identifies these systems through the **Managed By** column in the list of protection storage systems.

Some roles do not allow you to view the **Infrastructure > Storage** page to identify the relationships between systems and system pools. If your role does not allow you to view this information, coordinate storage target assignments with your system administrator.

Protection policies that target a system pool can replicate to another system pool or to a stand-alone protection storage system. Conversely, policies that target a stand-alone protection storage system can replicate to another protection storage system or to a system pool.

System pool reporting

Protection storage reporting differs slightly between individual protection storage systems and system pools. These differences are visible on the **Storage** page and the protection storage details pane.

The following table describes how specific columns in the list of protection storage systems behave for system pools.

Table 22. System pool reporting

Column	Description
Total	The total capacity of the system pool.
Available	The largest available space for storage unit placement on a single system in the system pool.
Free	The remaining unused space in the pool.
Encryption	On if any DD system in the system pool has enabled encryption.

Adding the values for **Available** and **Free** yields the total amount of unused space within the system pool.

Mobile DD Boost users

Smart Scale mobile DD Boost users own mobile storage units on system pools. This concept extends the association between DD Boost users and ordinary storage units to the system pool scope.

Mobile DD Boost users provide a unique user ID within a DDMC data center and control access to the associated mobile storage units. These users are centrally managed and unique across data centers.

Mobile DD Boost users send their requests to the DDMC instance which manages the entire system pool. DDMC, in turn, forwards the request to the correct system within the system pool.

As with other storage units, PowerProtect Data Manager associates a mobile DD Boost user with each mobile storage unit under the control of PowerProtect Data Manager.

Storage units provides more information about mobile storage units.

System pool limitations

Before you use system pools, review the following information:

- Storage Direct policies do not support system pools. The storage target list shows all protection storage systems, regardless of membership, but does not show system pools.
- Block volume policies do not support system pools.
- Cloud Tiering does not support system pools. If the primary backup or retention targets a system pool, you cannot add a Cloud Tiering objective to the protection policy. If the replication objective targets a system pool, you cannot add a Cloud Tiering objective to the replication objective.
- Server disaster recovery (DR) does not support system pools for protection policies. Protection policies that target system pools do not synchronize to the remote server.
- Server DR does not support system pools as a recovery target. The list of target protection storage systems does not include system pools.
- When automatically creating a mobile storage unit on a system pool for a protection policy:
 - If the policy encryption setting is enabled, PowerProtect Data Manager requests placement on a pool member where DD Boost file replication encryption is enabled.
 - If the policy encryption setting is disabled, PowerProtect Data Manager makes no specific placement request. The mobile storage unit may reside on a pool member where DD Boost file replication encryption could be either enabled or disabled.
 - The retention lock setting for the system pool and pool members must match the retention lock setting for the protection policy. If retention lock is disabled for the system pool or pool members but enabled for the protection policy, or conversely, mobile storage unit creation fails.

Mobile storage unit migration within a system pool

Review the following PowerProtect Data Manager prerequisites and postrequisites before you migrate mobile storage units within a system pool through DDMC.

During a migration, the selected storage units are unavailable for protection workflows. However, you can coordinate the backup and migration schedules to reduce downtime for the affected workflows.

You can only migrate to a destination that matches the requirements of the mobile storage unit. The PowerProtect DD documentation provides more information about these requirements.

Supported asset types

- VMware virtual machines
- Oracle databases
- Microsoft SQL Server databases
- Microsoft Exchange Server databases
- File systems
- Network-attached storage (NAS) shares
- SAP HANA databases
- Kubernetes clusters

Migration

Perform the following actions:

1. Review the PowerProtect DD documentation for migration instructions.
2. Start the migration and complete all steps leading up to the commit stage.
3. Before you commit the migration, stop the related PowerProtect Data Manager operations for the selected storage units. Stop PowerProtect Data Manager operations before mobile storage unit migration provides instructions.
4. Commit the migration and wait for migration to complete.
5. Restore full PowerProtect Data Manager operation. Restore PowerProtect Data Manager operations after mobile storage unit migration provides information.
6. Optionally, verify operation. Verify operation after mobile storage unit migration provides information.

Stop PowerProtect Data Manager operations before mobile storage unit migration

To quiesce PowerProtect Data Manager before you commit the migration, complete the following actions:

Steps

1. Disable any protection policies that use the selected storage units. Disable a protection policy provides instructions.
2. If the affected protection policies have replication objectives, perform manual replication to eliminate any replication backlog. Manual replication of protected assets provides instructions. Scheduled replication activities continue after you disable a protection policy.
3. Allow all running protection and restore activities for the affected protection policies to complete.
4. Disable server disaster recovery (DR). Disable server DR backups provides instructions.
5. Delete any Instant Access sessions which were started from the selected storage units. The *PowerProtect Data Manager Virtual Machine User Guide* provides instructions.
6. Disable compliance verification. The *PowerProtect Data Manager Security Configuration Guide* provides instructions.

Next steps

Refrain from the following activities until the migration completes and you resume normal operations:

- Performing manual backups of assets for the affected protection policies.
- Changing retention periods on the affected protection policies.

Restore PowerProtect Data Manager operations after mobile storage unit migration

To unquiesce PowerProtect Data Manager after migration, complete the following actions:

Steps

1. Enable compliance verification. The *PowerProtect Data Manager Security Configuration Guide* provides instructions.
2. Enable server DR. Manually configure server DR backups provides instructions.
3. Enable any protection policies that use the selected storage units. Enable a disabled protection policy provides instructions.

Verify operation after mobile storage unit migration

After you unquiesce PowerProtect Data Manager following a migration, optionally verify the operation of all protection policies that use the selected mobile storage units:

Steps

1. Perform a manual backup for each affected protection policy. Manual backups of protected assets provides instructions.
2. If the affected protection policies have replication objectives, perform manual replication. Manual replication of protected assets provides instructions.
3. Browse the existing and new backups of assets for the affected protection policies.
4. Verify that you can restore from the new backups and their replicas, including Instant Access restores.
5. Verify that you can delete existing backups and replicas. Delete backup copies provides instructions.

Add protection storage

Add and configure a storage system to use as a target for protection policies. Only the Administrator role can add protection storage.

Prerequisites

NOTE:

When adding a High Availability PowerProtect DD system, observe the following points:

- Do not add the individual active and standby DD systems to PowerProtect Data Manager.
- In the **Address** field, use the hostname that corresponds to the floating IP address of the High Availability PowerProtect DD system.
- The High Availability PowerProtect DD system is verified with the root certificate.

Steps

1. From the left navigation pane, select **Infrastructure > Storage**.
The **Storage** window appears.
2. In the **Protection Storage** tab, click **Add**.
3. In the **Add Storage** dialog box, select a storage system (**PowerProtect DD System** or **PowerProtect DD Management Center**).
For a system pool, select **DDMC**.
4. To add a High Availability PowerProtect DD system, select the check box.
5. Specify the storage system attributes:
 - a. In the **Name** field, specify a storage name.
 - b. In the **Address** field, specify the hostname, fully qualified domain name (FQDN), or the IP address.
 - c. In the **Port** field, specify the port for SSL communication. Default is 3009.
6. Under **Host Credentials** click **Add**. If you have already configured protection storage credentials that are common across storage systems, select an existing password. Alternatively, you can add new credentials, and then click **Save**.
7. If a trusted certificate does not exist on the storage system, a dialog box appears requesting certificate approval. Click **Verify** to review the certificate, and then click **Accept**.
8. Click **Save** to exit the **Add Storage** dialog and initiate the discovery of the storage system.
A dialog box appears to indicate that the request to add storage has been initiated.
9. In the **Storage** window, click **Discover** to refresh the window with any newly discovered storage systems.
When a discovery completes successfully, the **Status** column updates to **OK**. If **DDMC** is selected, all DD systems managed by the host will be listed after discovery.
10. To modify a storage system location, complete the following steps:
A storage system location is a label that is applied to a storage system. If you want to store your copies in a specific location, the label helps you select the correct storage system during policy creation.
 - a. In the **Storage** window, select the storage system from the table.
 - b. Click **More Actions > Set Location**.
The **Set Location** window appears.
 - c. Click **Add** in the **Location** list.
The **Add Location** window appears.
 - d. In the **Name** field, type a location name for the asset, and click **Save**.

Results

PowerProtect Data Manager displays the available protection storage systems. For each protection storage system, the **Managed By** column contains one of the following:

Table 23. Managed By column values

Protection storage type	Value
A stand-alone protection storage system.	The name of the protection storage system.

Table 23. Managed By column values (continued)

Protection storage type	Value
A protection storage system or a system pool that is managed by DDMC.	The name of the DDMC instance.

Edit protection storage

You can change the name, address, port number, and credentials for an existing protection storage system. Only the Administrator role can edit protection storage.

Prerequisites

Review the prerequisites for Server DR and Search Engine nodes. Change the IP address or hostname of a DD system provides more information. Review the limitations for each enabled asset source, as some asset sources may not support changing the address for a protection storage system.

Ensure that backup, restore, and FLR jobs are not running.

About this task

This task changes the stored Management interface for the protection storage system. If you change any other network interfaces, you must also update the preferred network interface for each protection policy objective that targets this protection storage system.

Steps


- From the left navigation pane, select **Infrastructure > Storage**.
The **Storage** window appears.
- In the **Protection Storage** tab, select a protection storage system and then click the link in the **Managed By** column.
The **Edit Storage** dialog box appears.
- In the **Edit Storage** dialog box, specify the storage system attributes:
 - In the **Name** field, specify a new storage name.
 - In the **Address** field, specify the new fully qualified domain name (FQDN) or the IP address.
 - In the **Port** field, specify the port for SSL communication. Default is 3009.
 - Under **Host Credentials**, select a new set of credentials or click **Add**.
- If a trusted certificate does not exist for the protection storage system, a dialog box appears requesting certificate approval. Click **Verify** to review the certificate, and then click **Accept**.
- Click **Save** to exit the **Edit Storage** dialog box.
- For Microsoft SQL Server or Oracle protection policies that use this protection storage system, update the lockbox:
 - From the left navigation pane, select **Protection > Protection Policies**.
The **Protection Policies** window appears.
 - Select the Microsoft SQL Server and Oracle protection policies that target this protection storage system, and then click **Set Lockbox**.

Replace protection storage

After you replace a DD protection storage system, you can use the PowerProtect Data Manager UI to update the required storage settings in PowerProtect Data Manager for the replacement storage. Only the Administrator role can update the protection storage settings in PowerProtect Data Manager. This process only applies to stand-alone DD systems (DD systems not managed by Dell PowerProtect DD Management Center).

The replacement of a DD protection storage system that is used with PowerProtect Data Manager can involve either of the following use cases:

- You perform a DD controller upgrade or head swap that replaces the DD system without requiring the migration of data.
- You perform collection replication or DD Cloud migration to migrate data from the original DD system to a new replacement system with the same or higher capacity.

 **CAUTION:** The new replacement DD must have all the storage units, DD Boost users, and data that existed on the original DD.

After you replace the DD system, ensure that you meet the following requirements before using the PowerProtect Data Manager UI to update the storage settings:

- All the protection policies that use the DD storage system are disabled.
- All the running protection jobs are completed.
- All the required data has been manually migrated as needed from the original DD to the replacement DD.

Use one of the following procedures to update the protection storage settings in PowerProtect Data Manager after replacing the DD storage system.

Updating the storage settings when data IPs are unchanged

Use the following procedure if you replaced the DD storage system and none of the network settings or data IPs have changed.

1. From the left navigation pane, select **Infrastructure > Storage**.


The **Storage** window appears.

2. On the **Protection Storage** tab, select the DD and then click **Discover**.

PowerProtect Data Manager automatically discovers the replacement DD storage and updates the storage settings for the replacement system. After the discovery, the replacement DD attributes appear on the **Protection Storage** tab of the **Infrastructure > Storage** window.

Updating the storage settings when data IPs are changed

Use the following procedure if you replaced the DD storage system and any of the network settings or data IPs have changed.

 **NOTE:** When data IPs have changed, update the network interface information for the replacement DD system.

1. From the left navigation pane, select **Infrastructure > Storage**.

The **Storage** window appears.

2. On the **Protection Storage** tab, select the DD and then select **More Actions > Replace System**.

3. Review the warning message about policies and running jobs, and then click **Continue**.


4. In the **Replace System** dialog box, specify the replacement DD system settings and then click **Apply**:

- In the **Address** field, specify the DD system management address as the fully qualified domain name (FQDN) or the IP address.
- In the **Port** field, specify the port for SSL communication. Default is 3009.
- Under **Host Credentials**, select a new set of credentials or click **Add**.

If a trusted certificate does not exist for the storage system, a dialog box appears requesting certificate approval. Click **Verify** to review the certificate, and then click **Accept**.

- Under **Network Mapping**, specify the replacement data IP for each listed data IP that has changed.

A system job is created for updating the storage settings for the replacement DD. You can monitor the system job in the **System Jobs** window.

 **NOTE:** You must wait until the replacement system job completes before you perform any PowerProtect Data Manager operations that use the DD system being replaced.

After the system job completes, the replacement DD attributes appear on the **Protection Storage** tab of the **Infrastructure > Storage** window. The protection policies that use the DD system are automatically updated with the replacement DD settings.

Storage units

PowerProtect Data Manager can create, configure, and reuse storage units on a protection storage system. These storage units are the targets for protection and replication policies.

The term "storage unit under the control of PowerProtect Data Manager" describes a storage unit that was created through one of the methods that are discussed here.

Review the applicable limitations before you create or change a storage unit, or change the protection or replication target for a policy. The *PowerProtect DD Virtual Edition Installation and Administration Guide* for the appropriate platform provides more information about storage units (MTrees).

Mobile storage units

For Smart Scale, mobile storage units extend the concept of a storage unit to the scope of an entire system pool. A mobile storage unit has the potential to move from one pool member to another. Thus:

- When you browse the storage units in a system pool, PowerProtect Data Manager displays only mobile storage units.
- When you browse the storage units on a DD system, PowerProtect Data Manager displays only regular (non-mobile) storage units.
- You must work with mobile storage units at the system pool level.


Aside from scope differences, PowerProtect Data Manager treats mobile storage units and regular storage units as equivalent.

Storage unit creation and configuration

PowerProtect Data Manager provides two ways to create storage units on the protection storage system:

- If you do not select an existing storage unit when you create a protection policy, PowerProtect Data Manager automatically creates a storage unit for you.
- Through the PowerProtect Data Manager UI, you can directly create storage units as required.

You can use the UI to configure the quotas and credentials for storage units under the control of PowerProtect Data Manager.

Click  to open the **Details** pane and see more information about an existing storage unit, including configuration values.

Storage unit selection

When you create or edit a protection policy, PowerProtect Data Manager provides the option to select a storage unit as the protection or replication target. The storage unit can be on the same or another protection storage system.

The **Storage** page lists all storage units that were discovered on a protection storage system. Only storage units created directly by PowerProtect Data Manager are available to select for a protection policy. Other storage units are not available to select, even if known.

A storage unit under the control of PowerProtect Data Manager can be the target for multiple protection policies. When you select an existing storage unit as a policy target, the policy inherits the storage unit's quota settings.

The instructions for creating and managing protection policies for each asset type provide more information about using storage units with policies.

Security

All protection policies and applications that share a storage unit can access any data in that storage unit. Reuse a storage unit only for policies and applications that belong to the same organizational unit or which share a trusted relationship. Policies and applications for different organizational units should use different storage units.

Any other external applications that also use the storage unit should protect and restrict access to the DD Boost credentials. These credentials provide access to the PowerProtect Data Manager data.

Automatic storage unit maintenance

For automatically-created storage units, automatic maintenance removes the storage unit when both the following conditions are true:

- No protection policies target the storage unit for backups or replication.
- The storage unit contains no backups.

Automatic maintenance removes these empty, unused storage units. For governance mode retention, automatic maintenance removes these storage units even if retention lock is enabled. Because deleting a storage unit with compliance mode enabled requires security officer credentials, automatic maintenance cannot remove these storage units.

For directly-created storage units, automatic maintenance does not remove the storage unit even when these conditions are true. In this case, contact the protection storage system administrator to remove the storage units.

Updating from previous releases

Any protection policy can use storage units that were automatically created for policies in a previous release of PowerProtect Data Manager. Policies that were created in a previous release continue to function as before.

Previous releases of the Oracle agent do not support storage units with multiple protection policies. The *PowerProtect Data Manager Oracle RMAN User Guide* provides more information.

Storage unit limitations

When using storage units with multiple protection policies, the following limitations apply:

- PowerProtect Data Manager cannot target or configure storage units that were not created through PowerProtect Data Manager.
- PowerProtect Data Manager cannot target storage units that were configured elsewhere for Cloud Tiering.
- Moving a protection policy to another storage unit or protection storage system may require a full backup.
 - For virtual machines, file system backups, Kubernetes, and Microsoft Exchange Server backups, the next backup is automatically promoted to a full backup.
 - For Microsoft SQL Server, Oracle, and SAP HANA backups, complete a manual full backup of these assets with the new storage unit.
- Protection policies for Storage Data Management cannot share a storage unit with other protection policies.
- Retention lock on a storage unit is disabled if any protection policy on that storage unit has retention lock disabled.
- Previous releases of the Oracle agent do not support sharing a storage unit between protection policies. The *PowerProtect Data Manager Oracle RMAN User Guide* provides more information.

Storage unit considerations for PowerProtect DD

With respect to PowerProtect DD, storage units have certain restrictions and best practices. Be aware of the following considerations:

- In order to avoid synchronization issues with PowerProtect Data Manager, any storage units that PowerProtect Data Manager is managing or using should not be modified directly from the DD.
- Storage units that you create in PowerProtect Data Manager must not be changed by the DD administrator to set up storage unit replication.
- Storage units that you create in PowerProtect Data Manager must not be configured for Cloud Tiering.
- For limitations that apply to supported storage units by PowerProtect DD model, see the E-Lab Navigator.

Retention locking

Retention locking prevents the deletion or alteration of data on a protection storage system for a specified period. PowerProtect Data Manager supports both governance mode and compliance mode retention locking for backups and replicas.

The PowerProtect DD documentation provides more information about each retention lock mode, including the differences between modes. Retention locking requires enablement and licensing on the protection storage system before use with PowerProtect Data Manager.

Retention locking is a two-stage process:

1. Create a storage unit on which you configure the appropriate retention lock mode. Configuration enables but does not activate retention locking.
2. Configure protection policies that both target this storage unit and activate retention locking. Toggling the retention lock setting for a protection policy activates retention locking in accordance with the configuration of the selected storage unit.

Once set, you cannot change the retention lock mode on a storage unit. To use a different retention lock mode with a protection policy, target a different storage unit. The original retention lock mode persists for existing backups or replicas that were created before the change.

The choice of retention lock mode may impact which protection policies can share a storage unit. Consider the retention lock settings when you design your storage unit architecture.

Compliance mode

Observe the following details before you configure or activate compliance mode retention locking:

- Compliance mode requires DDOS 7.10 or later. Earlier versions support only governance mode.
- Compliance mode requires the security officer credentials for the associated protection storage system. PowerProtect Data Manager does not store the security officer credentials.
- The Storage Direct agent for Storage Data Management does not support compliance mode.
- The option to create a storage unit through the selection drop-down list during protection policy configuration does not support compliance mode, only governance mode. To use compliance mode, create and configure a storage unit before you configure an associated protection policy.
- Deleting a storage unit with compliance mode enabled requires the security officer credentials for the associated protection storage system.

System pools and compliance mode retention locking

Mobile storage unit creation can place the storage unit on any pool member. However, the security officer credentials are unique to each pool member. Use the following roadmap to create a mobile storage unit and enable retention locking after creation.

1. Ensure that compliance mode is enabled for all pool members.
2. Create a mobile storage unit and set the retention lock mode to `None`.
3. Review the details for the mobile storage unit and note the pool member where the storage unit resides.
4. Edit the mobile storage unit and change the retention lock mode to compliance mode. Provide the security officer credentials for that pool member.

Create a storage unit

Directly create a storage unit through the PowerProtect Data Manager UI for use with protection policies.

Prerequisites

Add at least one protection storage system for PowerProtect Data Manager.

Steps

1. From the left navigation pane, select **Infrastructure > Storage**.
The **Storage** window appears.
2. On the **Protection Storage** tab, select a storage system, and then select **More Actions > Manage Storage Units**.
The **Storage Units** page opens and displays a list of the storage units under the control of PowerProtect Data Manager.
3. Select **Add**.
The **Create Storage Unit** or **Create Mobile Storage Unit** dialog box opens.
4. Type a name for the new storage unit.
5. For mobile storage units in system pools, select a **Network Group**.
Network groups are configured in DDMC to provide access to the system pool over different physical or virtual networks. A network group contains information about the IP addresses for the pool members and the IP address that clients use for access to the system pool.
6. Set the capacity and stream quotas that restrict the storage unit resource consumption.

There are two kinds of quota limits—hard limits and soft limits. You can set either a soft or hard limit or both a soft and hard limit. Both values must be integers, and the soft value must be less than the hard value.

i **NOTE:** When you set a soft limit and the limit is reached, an alert is generated, but data can still be written. When you set a hard limit and the limit is reached, data cannot be written. All data protection operations fail until data is deleted from the storage unit. The *PowerProtect DD Virtual Edition Installation and Administration Guide* for the appropriate platform provides more information about quota configuration.

- a. **Capacity Quota**—Controls the total size of precompression data that is written to the protection storage.
- b. **Stream Quota**—The number of concurrent streams allowed during data protection operations. Setting a **Stream Quota** limit can help ensure that performance is not impacted negatively when a data protection operation consumes too many resources.

7. Set a **Retention Lock Mode** from the available modes: *None*, *Compliance*, or *Governance*.

This field displays only the licensed and enabled options for the selected protection storage system. If no retention lock modes are enabled, the only option is *None*.

If you select *Compliance*, provide the username and password for the security officer who is associated with the protection storage system.

8. Select **Save**.

Results

PowerProtect Data Manager creates the storage unit on the selected protection storage system.

Edit a storage unit

Configure the settings for an existing storage unit through the PowerProtect Data Manager UI. You can also view a list of protection policies that target the storage unit.

About this task

Any changes to these storage unit attributes that you make directly on the protection storage system are also reflected in PowerProtect Data Manager.

Steps

1. From the left navigation pane, select **Infrastructure > Storage**.
The **Storage** window appears.
2. On the **Protection Storage** tab, select a storage system, and then select **More Actions > Manage Storage Units**.
The **Storage Units** page opens and displays a list of the storage units under the control of PowerProtect Data Manager.
3. To view the details or usage for a storage unit, select  for that storage unit.
The **Details** pane opens and displays the name, type, capacity, quota information, and a list of protection policies that currently target the storage unit.
The storage unit may contain copies from protection policies that no longer target the storage unit.
4. Select a storage unit from the list, and then select **Edit**.
The **Edit Storage Unit** or **Edit Mobile Storage Unit** dialog box opens.
5. For mobile storage units in system pools, select a **Network Group**.
Network groups are configured in DDMC to provide access to the system pool over different physical or virtual networks. A network group contains information about the IP addresses for the pool members and the IP address that clients use for access to the system pool.
6. Set the capacity and stream quotas that restrict the storage unit resource consumption.
There are two kinds of quota limits—hard limits and soft limits. You can set either a soft or hard limit or both a soft and hard limit. Both values must be integers, and the soft value must be less than the hard value.

i **NOTE:** When you set a soft limit and the limit is reached, an alert is generated, but data can still be written. When you set a hard limit and the limit is reached, data cannot be written. All data protection operations fail until data is deleted

from the storage unit. The *PowerProtect DD Virtual Edition Installation and Administration Guide* for the appropriate platform provides more information about quota configuration.

- a. **Capacity Quota**—Controls the total size of precompression data that is written to the protection storage.
 - b. **Stream Quota**—The number of concurrent streams allowed during data protection operations. Setting a **Stream Quota** limit can help ensure that performance is not impacted negatively when a data protection operation consumes too many resources.
7. If the **Retention Lock Mode** is *None*, set a **Retention Lock Mode** from the available modes: *Compliance* or *Governance*.

This field displays only the licensed and enabled options for the selected protection storage system. If no retention lock modes are enabled, the only option is *None*.

If you select *Compliance*, provide the username and password for the security officer who is associated with the protection storage system.

8. Select **Save**.

Results

PowerProtect Data Manager updates the storage unit settings.

Delete a storage unit

Because deleting a storage unit with compliance mode retention locking requires security officer credentials, automatic maintenance cannot remove these storage units. Instead, use this procedure to remove storage units with compliance mode retention locking.

Prerequisites

Before you can delete a storage unit, the storage unit must be empty and not targeted by any protection policies. The storage unit must be under the control of PowerProtect Data Manager and created by this instance of PowerProtect Data Manager.

If compliance mode retention locking is enabled, the security officer credentials for the associated protection storage system are required.

Steps

1. From the left navigation pane, select **Infrastructure > Storage**.
The **Storage** window appears.
2. On the **Protection Storage** tab, select a storage system, and then select **More Actions > Manage Storage Units**.
The **Storage Units** page opens and displays a list of the storage units under the control of PowerProtect Data Manager.
3. Select a storage unit from the list, and then select **Delete**.
The **Enter Security Officer Credential** dialog box opens.
4. Provide the security officer credentials and then click **OK**.

Results

PowerProtect Data Manager removes the storage unit.

Enable Indefinite Retention Hold on a storage unit

An indefinite retention hold (IRH) prevents any changes or deletion of data on a storage unit for an indefinite period. IRH prevents the disabling of Retention Lock Governance on a storage unit, while still allowing modification of retention lock attributes. In order for IRH to restrict the deletion or modification of files, the files must be locked or have expired locks. Files with no lock history remain unaffected by the IRH.

Prerequisites

Ensure that Retention Lock Mode (*Governance* or *Compliance*) is enabled on the storage unit.

Steps

1. From the left navigation pane, select **Infrastructure > Storage**.
The **Storage** window appears.
2. From the **Protection Storage** tab, select a storage system, and then select **More Actions > Manage Storage Units**.
3. Select a storage unit from the list, and then select **More Actions > Enable Indefinite Retention Hold**.
4. If required, provide the security officer user credentials, and then click **Enable**.
Security officer credentials are required when enabling or disabling IRH on a storage unit with Compliance mode enabled.

Disable Indefinite Retention Hold on a Storage Unit

Disable indefinite retention hold (IRH) on a storage unit to allow deletion of expired files or disabling of Retention Lock Governance on the storage unit.

Steps

1. From the left navigation pane, select **Infrastructure > Storage**.
The **Storage** window appears.
2. From the **Protection Storage** tab, select a storage system, and then select **More Actions > Manage Storage Units**.
3. Select a storage unit from the list, and then select **More Actions > Disable Indefinite Retention Hold**.
4. If required, provide the security officer user credentials, and then click **Disable**.
Security officer credentials are required when enabling or disabling IRH on a storage unit with Compliance mode enabled.

Working with storage unit passwords

The *PowerProtect Data Manager Security Configuration Guide* provides instructions for the following topics:

- Viewing an existing storage unit password
- Changing a storage unit password through the UI
- Changing the storage unit password policy

Differences in storage system and storage unit space reporting

Review the following sections for information about differences in the manner that storage space is reported in PowerProtect Data Manager.

Base 10 standard used for size calculations in the PowerProtect Data Manager UI

For size calculations (for example, asset size, storage system capacity), the PowerProtect Data Manager UI uses the Base 10 standard, which specifies the size in MB, GB, and TB.

Other components, however, might use the Base 2 standard, which specifies the size in MiB, GiB, and TiB. When there is a discrepancy in reported size, use the UI to obtain the most correct information.

How storage unit capacity is reported in PowerProtect Data Manager and DD Virtual Edition

Due to differences in space calculation (physical capacity vs logical capacity), there is a discrepancy between how storage unit capacity is reported in PowerProtect Data Manager and DD Virtual Edition.

For example, because PowerProtect Data Manager displays the DD storage unit logical capacity, the value that is reported when you select **More Actions > Manage Storage Units** in the PowerProtect Data Manager UI **Infrastructure > Storage** window might be greater than the amount reported in DDVE, which displays the physical capacity.

To determine the physical storage unit capacity, use DDVE instead.

Monitoring storage capacity thresholds

PowerProtect Data Manager periodically monitors protection storage usage and reports alerts when a system reaches two capacity thresholds. As a best practice, check for these alerts and respond before the system exhausts storage capacity.

At 80% capacity, PowerProtect Data Manager generates a weekly warning alert. At this threshold, you should develop a strategy to add capacity or move protection policies to another storage target. Managing Protection Policies provides more information about moving policies.

At 95% capacity, PowerProtect Data Manager generates a daily critical alert. At this threshold, capacity exhaustion is imminent.

Changing the capacity alerting thresholds requires contacting Support.

Using the PowerProtect Search Engine

Topics:

- PowerProtect Search Engine
- Set up and manage indexing
- Search Engine node deletion
- Edit the network configuration for a Search Engine node
- Perform a search
- Troubleshooting Search Engine issues

PowerProtect Search Engine

When you deploy PowerProtect Data Manager, the PowerProtect Search Engine software is installed by default.

The PowerProtect Search Engine indexes virtual machine file metadata to enable searches based on configurable parameters. To use this feature, add at least one Search Engine node to the Search Engine to form a cluster. Adding a Search Engine node enables the indexing feature.

You can enable the indexing option when creating protection policies so that the assets are indexed while they are backed up. Recovering indexes from a disaster is a manual process. The indexing recovery process will be automated in a future release.

When a DR backup is run, scheduled, or manually triggered, the cluster backup workflow backs up the cluster index data. A backup task is created, and you can view the individual status of the Search Component backup under **Details**.

NOTE: Scheduled backups with Search cluster integration appear in the Jobs pane as two identical jobs: an initialization job, which runs immediately, and the backup job, which runs both server DR and Search cluster backups.

Deploy Search Engine nodes with fully qualified domain names (FQDNs) only. PowerProtect Data Manager verifies that the hostname is an FQDN before deployment.

Limitations

- PowerProtect Search Engine is an optional feature that can be enabled, set up, and configured for virtual machine backups and protection policies. When you enable this feature, a backup of the Search Engine is taken as part of the server backup process. As of this release, you cannot disable these backups. Therefore, when **Search** is enabled, you must add the Search Engine node on the DD system that contains the ServerBackup MTree to the **Allow** list. If you use NFS for server DR, add the Search Engine node IP address or hostname to the client list for the NFS export.
- After an update to PowerProtect Data Manager, with the Search Engine already configured, and the first time that you use the **Networks** page to add a virtual network to an environment, PowerProtect Data Manager does not automatically add the virtual network to the Search Engine. Instead, manually edit each node to add the virtual network. This action makes the Search Engine aware of virtual networks. Any subsequent new virtual networks are automatically added to the Search Engine.
- If an operational Search Engine node fails, the node cannot be recovered, and the Search Engine cluster has a status of **Failed**, then the cluster must be removed and a new cluster created.

Set up and manage indexing

Set up a Search Engine node and configure indexing.

Prerequisites

Ensure that:

- A vCenter datastore has been configured. The *PowerProtect Data Manager Virtual Machine User Guide* provides detailed steps for adding a vCenter server as an asset source.
- PowerProtect Data Manager has discovered the networks for the vCenter server.
- The following requirements for the PowerProtect Search Engine are met:

NOTE: Each Search Engine node must meet the system requirements.

- CPU: 4 * 2 GHz (4 virtual sockets, 1 core for each socket)
 - Memory: 8 GB RAM
 - Disks: 3 disks (50 GB each) and 1 disk (1 TB)
 - Internet Protocol: Either only IPv4 or only IPv6
 - NIC: One vmxnet3 NIC with one port
- The PowerProtect Data Manager system is configured to use an NTP server. NTP server configuration is required to synchronize the time across the Search Engine nodes in a multi-node cluster.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Search Engine**, and then click **Add Node**.
2. In the **Add Search Engine Node** wizard, provide the required parameters.
 - **Hostname, IP Address, Gateway, DNS, and Netmask.**
 - **vCenter**—If you have added multiple vCenter server instances, select the vCenter server on which to deploy the Search Engine node.

NOTE: Ensure that you do not select the internal vCenter server.

- **ESX Host/Cluster**—Select on which cluster or ESXi host you want to deploy the Search Engine node.
- **Network**—Displays all the networks that are available under the selected ESXi Host/Cluster. For virtual networks (VLANs), this network carries management traffic.
- **Data Store**—Displays all datastores that are accessible to the selected ESXi Host/Cluster.

3. Click **Next**.
The **Networks Configuration** page displays.

4. On the **Networks Configuration** page:

The **Networks Configuration** page configures the virtual network (VLAN) to use for Data for Management Components traffic. To continue without virtual network configuration, leave the **Preferred Network Portgroup** selection blank and then click **Next**.

- a. From the **Preferred Network Portgroup** list, select a Virtual Guest Tagging (VGT) group.
VST (Virtual Switch Tagging) groups are not supported.

The list displays all virtual networks within the trunk range. If you select a portgroup that contains multiple networks, PowerProtect Data Manager automatically selects all networks. Individual networks cannot be selected.

A Search Engine node requires an IP address from the static IP pool for each selected virtual network. If there are not enough IP addresses in a pool, the wizard prompts you to supply additional addresses for that network.

Ensure that the selected virtual networks support a traffic type that is compatible with Search Engine nodes.

- b. If required, type an available static IP address or IP address range in the **Additional IP Addresses** column for the indicated virtual network.

For convenience when working with multiple virtual networks, you can also use one of the **Auto Expand** options:

- **Expand Last IP**—The wizard increments the host portion of the last IP address in the static IP pool. Click **Apply**.
- **Same Last Digit**—The wizard adds the network portion of the IP address to the specified value. Type the host portion of the IP address and then click **Apply**.

The wizard updates the value in the **Additional IP addresses** column for each network. Verify the proposed IP addresses.

- c. Click **Next**.
5. On the **Summary** page, review the information and then click **Finish**.
The new Search Engine node is deployed, and details are displayed in the lower panel.
 6. (Optional) Repeat the previous steps to deploy additional Search Engine nodes to the cluster.

NOTE: Ensure that the previous Search Engine node successfully deploys before you add another.

7. In the **Configure Search Engine** dialog box, enable or disable indexing, accept or change the expiration period, and then click **OK**.

NOTE:

- When the index cluster reaches 70 percent, an alert is generated. When it reaches 90 percent, an alert is generated and indexing is suspended. Specify a global index expiry interval to periodically clean up indexes, which frees up space.
- To turn off or modify indexing, select **Infrastructure > Search Engine**, select the cluster, and click **Configure Cluster**. From the **Configure Search Cluster** dialog box, you can enable/disable the service or change the number of expiration days.
- Indexes expire according to the global setting or when the associated copies expire, whichever occurs first.
- To stop indexing assets that have been added to a protected protection policy, disable the indexing option during protection policy configuration.
- You can add up to a maximum of 5 Search Engine nodes.

Next steps

NOTE:

When you edit or retry an operation that failed and there are additional IP addresses in the address pool, PowerProtect Data Manager marks the last failed IP address as abandoned. PowerProtect Data Manager does not try to reuse any IP addresses that are marked as abandoned. The UI does not display this condition.

KB article 000181120 provides more information about how to use the REST API to detect when an IP address is marked as abandoned. The article also provides steps to correct this condition so that the IP address can be used again.

Search Engine node deletion

PowerProtect Data Manager UI supports the deletion of a Search Engine node from a multinode cluster.

The following remediations can be performed on a Search Engine node:

- Delete an operational node from a Search Engine cluster to decrease the cluster capacity when the space is no longer required.
- Redeploy or delete a node that failed to deploy to a Search Engine cluster.
- Delete all nodes in a Search Engine cluster to remove the cluster.

Delete an operational Search Engine node

You can delete an operational Search Engine node to decrease the cluster capacity when the space is no longer required.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Search Engine**.
2. Select the node from the list that you want to delete, and then select **More Actions > Delete Node**.
3. In the **Delete Search Engine Node** window, click **Delete Node**.

CAUTION: Do not select **Delete node without moving the index data**. If you select this option, the Search Engine cluster becomes inactive and cannot be recovered.

4. Go to the **Jobs > System Jobs** window to monitor the progress of the node deletion.

Redeploy or delete a Search Engine node that failed to deploy

PowerProtect Data Manager enables you to redeploy or delete a Search Engine node that could not be successfully deployed.

About this task

The **Redeploy Node** functionality is only enabled for nodes that have been added but could not be successfully deployed to the Search Engine.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Search Engine**.
2. Select the node that failed to deploy.
3. Do one of the following:

- To redeploy the node, select **More Actions > Redeploy Node**.

The **Redeploy Search Engine Node** wizard opens. The Search Engine populates the fields with the information that you supplied when you added the node. Verify that the information is correct.

- To delete the node, select **More Actions > Delete Node**.
4. Go to the **Jobs > System Jobs** window to monitor the progress of the node redeployment or deletion.

Delete all Search Engine nodes to remove the Search Engine cluster

You can delete all Search Engine nodes in a Search Engine cluster to remove the cluster.

About this task

Removing the cluster is necessary if one of the operational nodes has failed, the node cannot be recovered, and the cluster has a status of **Failed**.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Search Engine**.
2. Perform the following steps for each of the nodes:

- a. Select the node from the list.
- b. Select **More Actions > Delete Node**.
- c. Click **Delete Node**.

NOTE: If no operational nodes have failed, the option **Delete nodes without moving the index data** is available. Selecting this option results in the cluster becoming irrecoverably inactive.

- d. Go to the **Jobs > System Jobs** window to monitor the progress of the node deletion operation.

Results

All Search Engine nodes are deleted and the Search Engine cluster is removed.

Edit the network configuration for a Search Engine node

To change the virtual network configuration, perform the following steps. To change any other network configuration settings, contact Customer Support.

Prerequisites

Before you remove a network, disable indexing. Set up and manage indexing provides instructions.

About this task

If Search Engine node deployment failed because of a virtual network configuration problem, you can update the configuration to add additional IP addresses to the static IP pool. If you did not configure a virtual network during initial deployment, you can also add the Search Engine node to a virtual network in the same Virtual Guest Tagging (VGT) port group.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Search Engine** and then select the applicable Search Engine node.
2. Select **More Actions > Edit Networks**.
The **Edit Search Engine Node** wizard opens to the **Network Configuration** page.
3. If applicable, from the **Preferred Network Portgroup** list, select a VGT network to carry Data for Management Components traffic.

The list displays all virtual networks within the trunk range. If you select a portgroup that contains multiple networks, PowerProtect Data Manager automatically selects all networks. Individual networks cannot be selected.

A Search Engine node requires an IP address from the static IP pool for each selected virtual network. If there are not enough IP addresses in a pool, the wizard prompts you to supply additional addresses for that network.

Virtual networks with a warning symbol (⚠) beside the network name require attention and review. For example, if you changed the network configuration, the configured traffic types may not support Search Engine nodes. Clear any interfaces which no longer apply to the Search Engine node.

4. If required, type an available static IP address or IP address range in the **Additional IP Addresses** column for the indicated virtual network.

For convenience when working with multiple virtual networks, you can also use one of the **Auto Expand** options:

- **Expand Last IP**—The wizard increments the host portion of the last IP address in the static IP pool. Click **Apply**.
- **Same Last Digit**—The wizard adds the network portion of the IP address to the specified value. Type the host portion of the IP address and then click **Apply**.

The wizard updates the value in the **Additional IP addresses** column for each network. Verify the proposed IP addresses.

5. Click **Next**.
6. On the **Summary** page, review the information and then click **Finish**.

Next steps

If you disabled indexing, re-enable indexing. Set up and manage indexing provides instructions.

Perform a search

When the Search Engine is deployed and configured, you can use the **File Search** functionality in the PowerProtect Data Manager UI to search across all indexed data to locate protected files and folders within backup copies. When asset types are set up for index searching, the **File Search** button appears in the **Restore** menu for assets.

Before performing a search, ensure that:

- A Search Engine node is set up.
- Search indexing is enabled.

Troubleshooting Search Engine issues

This section lists troubleshooting for Search Engine issues.

Some Search Engine troubleshooting procedures require the credentials for individual Search Engine nodes. Search Engine nodes have admin and root user accounts that are used for troubleshooting software issues. The *PowerProtect Data Manager Security Configuration Guide* provides instructions to manage Search Engine node credentials.

Error displays during Search Engine node failure

The following error might display during a search when a Search Engine node fails:

```
Not able to deploy search-node.com. Another session "<host_name>" is already configured with the same hostname. Would you like to redeploy search_node or delete the node?
```

If this error occurs, delete the Search Engine node, and then retry the operation. If you choose to edit, delete the node. The new mode modal then appears with your previous inputs. The input that caused the error is marked as critical.

Certificate issues

Issues with indexing backups and/or performing search queries might result when certificates that were deployed on the Search Engine node were corrupted.

Perform one of the following tests to determine certificate issues:

- Use the log bundle download utility in PowerProtect Data Manager to examine the Backup VM logs in VM Direct, and look for a log entry like the following:

```
ERROR: Failed to Upload File: /opt/emc/vproxy/runtime/tmp/vproxyd/plugin/search/e6c356a1-fbaf-4231-9f6f-a0166b74909a/<search node>-e081fdea-3599-4a6c-abc4-1b5487cb9a32-e523a94c-2d01-5234-ab3c-7771cfab3c58-7f16bcbb72d7b49ea073356f0d7388ac08461827.db.zip to https://<search node>:14251/upload, Error sending data chunk. Post https://<search node>:14251/upload: x509: certificate signed by unknown authority (possibly because of "crypto/rsa: verification error" while trying to verify candidate authority certificate "PPDM Root CA ID-d5ec56b8-69ec-4183-9c94-7c0230408765")
```

- Examine the REST engine logs in the Search Engine node (`/opt/emc/search/logs/rest-engine/*.log`), and look for certificate verification errors.
- Run a search either through the UI or through the API `<PowerProtect Data Manager>/api/v2/file-instances` and look for a certification verification error.

Examine the certificate files on each Search Engine node to investigate further. If necessary, regenerate the certificate files.

Verify certificates

Use this procedure to verify that certificates are valid and uncorrupted:

1. Verify that the `rootca.pem` file is the same in all the relevant nodes (Search Engine node, PowerProtect Data Manager, and VM Direct node).

NOTE: The `rootca.pem` file name is different on each node:

- PowerProtect Data Manager—`/etc/ssl/certificates/rootca/rootca.pem`
- Search Engine node—`/var/lib/dellemc/vmboot/trust/thumbprint`
- VM Direct—`/var/lib/dellemc/vmboot/trust/thumbprint`

2. Run the following OpenSSL command to find out whether the root certificate file is corrupt or invalid: `openssl verify <rootca.pem>`

Response:

```
/var/lib/dellemc/vmboot/trust/thumbprint: C = US,
O = DELL Corporation,
CN = PPDM Root CA ID-4c9de850-24ab-42ec-a9a7-6080849d0d24

error 16 at 0 depth lookup:self signed certificate
```

OK

Ensure that the CN values match.

Certificate verification fails

If the certificate verification steps fail, you must re-create the certificates on the Search Engine node or VM Direct node:

1. Connect to the PowerProtect Data Manager console and change to the root user.
2. Use the `Get` command in the `infranodemgmt` utility to determine the Search Engine node FQDN.
3. Run `/usr/local/brs/puppet/scripts/generate_certificates.sh -n -c -b <node FQDN>`
A properties file is created in the `/root` directory called `<node FQDN>.properties`.
4. Open this file to determine the location of the generated certificates. They should be located in `/etc/ssl/certificates/<node FQDN>`.
5. Obtain the Search Engine node credentials. The *PowerProtect Data Manager Security Configuration Guide* provides instructions.
6. From a separate terminal, SSH into the Search Engine node.
7. Change directory to `/var/lib/dellemc/vnboot/trust` and move the `key`, `cert`, and `thumbprint` files over.
8. Copy the certificate files that were generated in PowerProtect Data Manager as follows:
 - `rootca.pem` to `thumbprint`
 - `<search node FQDN>key.pem` to `key`
 - `<search node FQDN>.pem` to `cert`.
9. Paste the files to `/var/lib/dellemc/vnboot/trust`.
10. Set the permissions for the `key`, `cert`, and `thumbprint` files to `0644`, and then set the ownership of these files to `root:app`.
11. Restart the REST engine service to pick up the new certificates: `systemctl restart search-rest-engine`.
12. Check the REST engine log file (`/opt/emc/search/logs/rest-engine/rest-engine-daemon-<fqdn>.log`) to verify that the service started successfully.

Ensure that the following message appears:

```
A valid Root CA certificate of backup server was provided during deployment
```

Result: Backup with indexing executes successfully and the Search Engine is functional.

Search Engine cluster is full

If the Search Engine is full, you can deploy additional nodes by following the steps in *Set up and manage indexing*.

If the Search Engine runs out of space and you do not want to deploy an additional node, you have the following options:

- Disable the service
- Shorten the expiration time to remove indexes sooner
- Remove indexes manually

To disable the service, complete the following steps:

1. From the PowerProtect Data Manager UI, select **Infrastructure > Search Engine**.
2. Select the cluster, and then click **Configure Cluster**.
3. In the **Configure Search Cluster** dialog box, switch the **Search Indexing** button to turn it off, and then click **Save**.

NOTE: This setting applies to all indexes in all protection policies in the Search Cluster.

To shorten the expiration time to remove indexes sooner, complete the following steps:

1. From the PowerProtect Data Manager UI, select **Infrastructure > Search Engine**.
2. Select the cluster, and then click **Configure Cluster**.
3. In the **Configure Search Cluster** dialog box, modify the **Search Index Expiration** and click **Save**. A recommended formula to determine the expiration time is: `Delete Index when Today = Backup-Date + Expiration Days + 1 day`. That is, one day after the backup expires.

NOTE: This setting applies to all indexes in all protection policies in the Search Engine.

To remove indexes manually, complete the following steps:

1. Use SSH to log in to the Search Engine.

2. Create a snapshot of the cluster using the following format:

```

{
  Command: "APP_SNAPSHOT",
  Title: "Initiate Index/Search Cluster Snapshot Process",
  AsyncCmd: false,
  Properties: {
    "Name": {
      Description: "Used to uniquely identify a particular snapshot",
      Type: STRING
    },
    "Action": {
      Description: "Action to perform, 'Create', 'Delete', 'Restore' or 'Cancel' a Snapshot",
      Type: STRING
    },
    "NFSHost": {
      Description: "NFS Host serving snapshot backup area.",
      Type: STRING
    },
    "NFSExport": {
      Description: "NFS Export path to mount too.",
      Type: STRING
    },
    "NFSDirPath": {
      Description: "NFS directory path to write too.",
      Type: STRING
    }
  }
}

```

For example:

```

{
  "Command": "APP_SNAPSHOT",
  "Title": "",
  "AsyncCmd": false,
  "Properties": {
    "Action": {
      "Description": "",
      "Required": false,
      "Type": "string",
      "IsArray": false,
      "Value": "Create",
      "Default": null
    },
    "Name": {
      "Description": "",
      "Required": false,
      "Type": "string",
      "IsArray": false,
      "Value": "DataManager_Catalog_Cluster_snapshot_2019-10-16-12-57-16",
      "Default": null
    },
    "NFSHost": {
      "Value": "10.25.87.88"
    },
    "NFSExport": {
      "Value": "/mnt/shared"
    },
    "NFSDirPath": {
      "Value": ""
    }
  }
}

```

3. You can delete indexes by protection policy or by asset. If the JSON command is stored at /home/admin/remove-plc.json, run the command, ./searchgmt -I /home/admin/remove-plc.json.

- Use the following format to delete indexes by protection policy:

```
{
  "Command": "APP_REMOVE_ITEMS",
  "AsyncCmd": false,
  "Properties": {
    "Action": {
      "Description": "Action to perform,
'AssetDelete', 'PLCDelete'",
      "Required": true,
      "Value": "PLCDelete",
    },
    "PLCID": {
      "Description": "PLC ID of item(s) to delete.",
      "Required": true,
      "Value": "7676d753-b57e-a572-6daf-33689933456d",
    }
  }
}
```

- Use the following format to delete indexes by asset type:

```
{
  "Command": "APP_REMOVE_ITEMS",
  "AsyncCmd": false,
  "Properties": {
    "Action": {
      "Description": "Action to perform,
'AssetDelete', 'PLCDelete'",
      "Required": true,
      "Value": "AssetDelete",
    },
    "AssetID": {
      "Description": "Optional, Asset ID of item(s)
to delete.",
      "Required": false,
      "Value": "503dd753-b57e-a572-6daf-44680033755f",
    },
    "PLCID": {
      "Description": "PLC ID of item(s) to delete.",
      "Required": true,
      "Value": "7676d753-b57e-a572-6daf-33689933456d",
    }
  }
}
```

NOTE:

- The time to complete the execution of these procedures depends on the number of backup copy asset indexes being deleted.
- This procedure does not impact regular operation of the cluster.

Troubleshooting a locked Search Engine node

The *PowerProtect Data Manager Security Configuration Guide* provides information about Search Engine node user accounts and credentials, including password management policies. The password management policies for these accounts are set to lock the admin user account after three failed attempts within five minutes. If you try to access the node while the admin user account is locked, the amount of time that the account remains locked increases.

A Search Engine node might become locked for the following reasons:

- A user or program makes three failed attempts to SSH into the Search Engine node.
- Running monitoring software that tries to log in to the Search Engine node with the wrong admin credentials.
- Running penetration testing on the virtual machines in a vCenter server.

The Search Engine node admin user accounts enable PowerProtect Data Manager to perform operations on each node, such as obtaining the health status of the node. If the account is locked, the health status of the node is reported as "Failed." When one of the nodes in the cluster is in a failed state, the entire cluster becomes unavailable. As a result, the cluster is unable to perform any indexing or search operations.

Workaround

To work around this issue, reset the Search Engine node admin credentials. Before you reset the credentials, determine why the admin account is locked.

Obtain the Search Engine node root credentials. Then, reset the Search Engine node admin credentials. The *PowerProtect Data Manager Security Configuration Guide* provides instructions.

Managing Assets

Topics:

- About asset sources, assets, and storage
- About other asset sources
- Prerequisites for discovering asset sources
- Enable an asset source
- Delete an asset source
- Adding a Cloud Snapshot Manager tenant

About asset sources, assets, and storage

In PowerProtect Data Manager, assets are the basic units that PowerProtect Data Manager protects. Asset sources are the mechanism that PowerProtect Data Manager uses to manage assets and communicate with the protection storage where backup copies of the assets are stored.

PowerProtect Data Manager supports PowerProtect DD Management Center (DDMC) as the storage and programmatic interface for controlling protection storage systems.

Supported asset sources

Asset sources can be a vCenter server, Kubernetes cluster, Network Attached Storage (NAS) appliance or share, application host, SMIS server, storage array, or Cloud Snapshot Manager tenant. Assets can be virtual machines, Kubernetes namespaces and persistent volume claims (PVCs), NAS appliance or share assets, Microsoft Exchange Server databases, Microsoft SQL Server databases, Oracle databases, SAP HANA databases, file systems, storage groups, or PowerStore block volumes.

Supported asset languages

PowerProtect Data Manager supports the protection of data hosted on operating systems in multiple languages. For more information, see the E-Lab Navigator.

Prerequisite to adding an asset source

Before you can add an asset source, you must enable the source within the PowerProtect Data Manager user interface.


In the **Assets** window, you can export asset records by using the **Export All** functionality.

IPv6 information not displayed by the **Asset Sources** window

The **Asset Sources** window does not display IPv6 information. If an asset only uses IPv6, the **IPv4** column displays a blank entry. To select an IPv6-only asset, refer to the **Name** column.

Maximum supported number of characters in an asset or storage name is 25

PowerProtect Data Manager does not support more than 25 characters in an asset or storage name.

 **CAUTION: If this maximum is exceeded, protection policy configuration fails.**

About other asset sources

In addition to vCenter server asset sources, PowerProtect Data Manager provides the option to enable other asset sources to protect additional asset types.

The *PowerProtect Data Manager Administrator Guide* does not provide instructions for Kubernetes clusters or agent asset-source management. Refer to the PowerProtect Data Manager online help or individual Kubernetes and agent user guides for more information.

NOTE: When following an agent user guide to install an agent, ensure that the drive or partition of the installation directory has sufficient free space.

The following other asset sources are supported:

File System agent

After the File System agent is approved and registered in the PowerProtect Data Manager UI, PowerProtect Data Manager integrates with the agent to enable an application administrator to protect and recover data on the File System host, and to check and monitor backup compliance against protection policies.

Kubernetes cluster

After the Kubernetes cluster asset source is added and registered in the PowerProtect Data Manager UI, PowerProtect Data Manager enables protection of PVCs and namespace data on the Kubernetes or Tanzu Kubernetes cluster.

NAS agent

After the NAS asset source is added and registered in the PowerProtect Data Manager UI, PowerProtect Data Manager enables protection of NAS assets.

Microsoft application agent for Microsoft Exchange Server

After the Microsoft application agent is approved and registered in the PowerProtect Data Manager UI, PowerProtect Data Manager integrates with the agent to enable an application administrator to protect and recover the Microsoft Exchange Server application data on the application host, and to check and monitor backup compliance against protection policies.

Microsoft application agent for Microsoft SQL Server

After the Microsoft application agent is approved and registered in the PowerProtect Data Manager UI, PowerProtect Data Manager integrates with the agent to enable an application administrator to protect and recover the Microsoft SQL Server application data on the application host, and to check and monitor backup compliance against protection policies.

Oracle RMAN agent

After the Oracle RMAN agent is approved and registered in the PowerProtect Data Manager UI, PowerProtect Data Manager integrates with the agent to enable an application administrator to protect and recover the Oracle application data on the application host, and to check and monitor backup compliance against protection policies.

SAP HANA agent

After the SAP HANA agent is approved and registered in the PowerProtect Data Manager UI, PowerProtect Data Manager integrates with the agent to enable an application administrator to protect and recover the SAP HANA application data on the application host, and to check and monitor backup compliance against protection policies.

Storage Direct agent for Storage Data Management

Storage Data Management uses snapshot backup technology to protect data on VMAX and PowerMax storage arrays by moving storage group data from the array to a DD system. After the Storage Direct agent is approved and registered in the PowerProtect Data Manager UI, and the DD system and the SMIS server are added and discovered, the Storage Direct agent enables you to discover the storage groups in the storage arrays, and assign unprotected storage groups to a protection policy for backup and recovery operations.

Storage Arrays

For the integration with Dell PowerStore, PowerProtect Data Manager provides centralized backup and restore operations to protect data on storage arrays. After the storage array asset source is added and discovered in the PowerProtect Data Manager UI, PowerProtect Data Manager enables protection of block volume assets.

Prerequisites for discovering asset sources

Review the following requirements before discovering an asset source.

- Ensure that the PowerProtect Data Manager is deployed and configured in the environment. The PowerProtect Data Manager deployment guides provide information.
 - Log in as a user with the Administrator role. Only the Administrator role can manage asset sources.
 - For a new system, enable one or more asset sources for the types of assets that you want to protect. Enable an asset source provides more information.
 - Configure all asset sources with an NTP server.
 - Remove any managed snapshots from virtual machines that will be configured to use the Transparent Snapshots Data Mover (TSDM) protection mechanism.
 - Before you register a Microsoft SQL Server application, ensure that the DD system has been discovered successfully.
 - For discovery of application agents and File System asset sources:
 - Ensure that all clocks on the application and File System hosts and PowerProtect Data Manager are time-synchronized to the local NTP server to ensure discovery of the backups.
 - Ensure that the application and File System hosts and the PowerProtect Data Manager network can see and resolve each other.
 - Ensure that port 7,000 is open on the application and File System hosts.
 - Discovery of a vCenter Server asset source excludes the following:
 - Virtual machines with a status of **Inaccessible**, **Invalid**, or **Orphaned**.
 - The virtual machine template.
 - The shadow or standby virtual machine created by RecoverPoint for Virtual Machines, also referred to as the vRPA copy.
 - The vSphere Cluster Service (vCLS) virtual machine.
- NOTE:** Virtual machines created by the vCLS are managed by VMware, and do not require PowerProtect Data Manager protection. Even when selected as part of a container, they are automatically excluded from protection. The `vmdm-discovery.log` provides a list of vCLS virtual machines that are excluded from protection.

Prior to performing the vCenter discovery, verify the status of any virtual machines that you want to discover.

Discovering asset sources in an opaque network

PowerProtect Data Manager supports the discovery of vCenter servers that are located in an opaque network.

VMware considers a network to be opaque if a vCenter server located in it is not managed by NSX or vSphere. vCenter servers on opaque networks can be discovered and their assets protected in the same way as vCenter servers that are managed by NSX or vSphere.

Discovering asset sources in a GCVE environment

There are special discovery considerations in a GCVE environment. Discovery fails unless GCVE-located vCenter servers have additional permissions.

Ensure the following permissions of any GCVE-located vCenter server:

- The `GVE.LOCAL\CloudOwner` user is mapped to the `Cloud-Owner-Role` role at the vCenter level.
- The `GVE.LOCAL\CloudOwner` to `Cloud-Owner-Role` mapping is not restricted to a lower-level container object in the vSphere object hierarchy.

Full discovery of application asset sources

If some application assets are not discovered, you can perform an immediate full discovery of application asset sources by using the on-demand discovery feature in the PowerProtect Data Manager UI.

Full discovery is available for the following application asset sources:

- Microsoft SQL Server
- Microsoft Exchange Server
- Oracle
- SAP HANA
- File System

To initiate a full discovery of application asset sources, complete the following steps:

1. Select **Infrastructure > Asset Sources**.
2. Select an application asset source and click **Discover**.
3. Select the **Initiate a full discovery** option, and then click **Yes**.

Enable an asset source

An asset source must be enabled in PowerProtect Data Manager before you can add and register the asset source for the protection of assets.

About this task

Only the Administrator role can manage asset sources.

In some circumstances, the enabling of multiple asset sources is required. For example, a vCenter Server and a Kubernetes cluster asset source must be enabled for Tanzu Kubernetes guest cluster protection.

There are other circumstances where enabling an asset source is not required, such as the following:

- For application agents and other agents such as File System and Storage Direct, an asset source is enabled automatically when you register and approve the agent host. For example, if you have not enabled an Oracle asset source but have registered the application host through the API or the PowerProtect Data Manager user interface, PowerProtect Data Manager automatically enables the Oracle asset source.
- When you update to the latest version of PowerProtect Data Manager from an earlier release, any asset sources that were previously enabled appear in the PowerProtect Data Manager user interface. On a new deployment, however, no asset sources are enabled by default.

Steps

1. From the PowerProtect Data Manager user interface, select **Infrastructure > Asset Sources**, and then click **+** to reveal the **New Asset Source** tab.
2. In the pane for the asset source that you want to add, click **Enable Source**.
The **Asset Sources** window updates to display a tab for the new asset source.

Results

You can now add or approve the asset source for use in PowerProtect Data Manager. For a vCenter server, Kubernetes cluster, SMIS Server, or PowerProtect Cloud Snapshot Manager tenant, select the appropriate tab in this window and click **Add**. For an application host, select **Infrastructure > Application Agents** and click **Add** or **Approve** as required.

NOTE: Although you can add a Cloud Snapshot Manager tenant to PowerProtect Data Manager in order to view its health, alerts, and the status of its protection, recovery, and system jobs, you cannot manage the protection of its assets from PowerProtect Data Manager. To manage the protection of its assets, use Cloud Snapshot Manager. For more information, see the *PowerProtect Cloud Snapshot Manager Online Help*.

Disable an asset source

If you enabled an asset source that you no longer require, and the host has not been registered in PowerProtect Data Manager, perform the following steps to disable the asset source.

About this task

NOTE: An asset source cannot be disabled when one or more sources are still registered or there are backup copies of the source assets. For example, if you registered a vCenter server and created policy backups for the vCenter Server virtual machines, then you cannot disable the vCenter Server asset source. But if you register a vCenter server and then delete it without creating any backups, you can disable the asset source.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Asset Sources**, and then select the tab of the asset source that you want to disable.
If no host registration is detected, a red **Disable** button appears.
2. Click **Disable**.

Results

PowerProtect Data Manager removes the tab for this asset source.

Delete an asset source

If you want to remove an asset source that you no longer require, perform the following steps to delete the asset source in the PowerProtect Data Manager UI.

About this task

Only the Administrator role can manage the asset sources.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Asset Sources**, and then select the tab for the type of asset source that you want to delete.
2. Select the asset source name in the asset source list, and then click **Delete**.
3. At the warning prompt that appears, click **Continue**.
The asset source is deleted from the list.

Results

PowerProtect Data Manager removes the specified asset source in the **Asset Sources** window.

For all asset sources except the vCenter Server, any associated assets that are protected by the protection policy are removed from the protection policy and their status is changed to **deleted**. These assets are removed automatically as part of daily PowerProtect Data Manager cleanup after all associated backup copies have been deleted. These assets can also be removed manually. The *PowerProtect Data Manager Administrator Guide* provides details on how to remove assets from PowerProtect Data Manager.

The copies of assets from the asset source are retained (not deleted). You can delete the copies from the copies page, if required.

Adding a Cloud Snapshot Manager tenant

After you enable the Cloud Snapshot Manager tenant asset-source with PowerProtect Data Manager, you use the **Asset Sources** window in PowerProtect Data Manager to add a Cloud Snapshot Manager tenant to the PowerProtect Data Manager environment.

Adding a Cloud Snapshot Manager tenant is required if you want to view Cloud Snapshot Manager jobs, alerts, and reports from a consolidated PowerProtect Data Manager dashboard.

Add a Cloud Snapshot Manager tenant

Perform the following steps to add a Cloud Snapshot Manager tenant as an asset source in the PowerProtect Data Manager UI.

Prerequisites

- Ensure that the asset source is enabled.
Enable an asset source provides instructions.
- Log in as a user with the Administrator role. Only the Administrator role can manage asset sources.
- The PowerProtect Data Manager server has Internet access and is able to reach <https://ssgosga.emc.com>.
NOTE: If this access is removed during normal operation, any existing Cloud Snapshot Manager information will continue to be displayed in the **Dashboard** window, but there will be no updates until Internet access is restored.
- This procedure requires the entry of values specific to Cloud Snapshot Manager. For more information, see the *PowerProtect Cloud Snapshot Manager Online Help*.

Steps

1. From the left navigation pane, select **Infrastructure > Asset Sources**.
The **Asset Sources** window appears.
2. Select the **Cloud Snapshot Manager** tab.
3. Click **Add**.
The **Add Cloud Snapshot Manager Account Details** dialog displays.
4. In the **Name** field, enter a descriptive name for the Cloud Snapshot Manager tenant.
5. In the **Tenant ID** field, enter the Cloud Snapshot Manager tenant ID.
6. Click the drop-down control next to **Cloud Snapshot Manager Credentials**, and then click **Add Credentials**.
 - a. In the **Name** field, enter the name of the Cloud Snapshot Manager tenant credentials.
 - b. In the **Client ID** field, enter the ID of the Cloud Snapshot Manager tenant.
 - c. In the **Client Secret** field, enter the secret of the Cloud Snapshot Manager tenant.
 - d. Click **Save**.
7. Click **Save**.

Results

PowerProtect Data Manager can view jobs, alerts, and reports related to protected Cloud Snapshot Manager resources.

Managing Protection Policies

Topics:

- Protection policies
- Before you create a protection policy
- Adding or editing a protection policy
- Viewing a summary of protection policies
- Run an asset-protection report
- Add a service-level agreement
- Run a compliance report
- Disable a protection policy
- Delete a protection policy
- Overview of PowerProtect Data Manager Cloud Tier
- Extended retention for protection policies created in PowerProtect Data Manager 19.11 or earlier
- Manual backups of protected assets
- Manual replication of protected assets
- Manual Cloud Tiering of protected assets
- Delete backup copies
- Removing expired backup copies
- Removing assets from PowerProtect Data Manager
- Protecting client assets after a client hostname change
- IfGroup configuration and PowerProtect Data Manager policies
- Troubleshooting failed replication jobs

Protection policies

Protection policies define sets of objectives that apply to specific periods of time. These objectives drive configuration, active protection, and copy-data-management operations that satisfy the business requirements for the specified data. Each policy type has its own set of user objectives.

Only the Administrator role can create or edit protection policies.

You can create protection policies for the following asset types:

- VMware virtual machines
- Microsoft Exchange Server databases
- Microsoft SQL Server databases
- Oracle databases
- SAP HANA databases
- File systems
- Kubernetes clusters
- Storage groups
- Block volumes
- Network-attached storage (NAS)

For each policy type, refer to the individual user guides.

In the **Protection Policies** window, you can export protection policy data by using the **Export All** functionality.

Before you create a protection policy

Consider the following best practices before creating a protection policy.

- You can only protect an asset with one policy at a time. Protection rules do not automatically move assets that were manually added to a policy to a different policy.
NOTE: If a Microsoft SQL Server is installed on a virtual machine, you can protect the Microsoft SQL Server database with an application-consistent backup without interfering with the Microsoft SQL Server agent-based backup.
- When creating a policy, limit the number of database assets within the policy to under 500 and stagger the start time of replication policies. These actions prevent potential replication failures.
- Before adding replication to a protection policy, ensure that you add remote protection storage as the replication location. Add protection storage provides detailed instructions about adding remote protection storage.
- Before scheduling weekly, monthly, or yearly backups, ensure that the PowerProtect Data Manager time zone is set to the local time zone.

Understanding backup technologies

PowerProtect Data Manager uses block-based backup technology when performing full or synthetic-full backups. The File System agent scans a volume or disk and backs up every block on the file system that is allocated to it. If only data that has changed is backed up, the block-based backup uses Changed Block Tracking.

Block-based backups support the following capabilities:

- High-performance backups with a predictable backup window
- Efficient backups of the deduplicated file systems used by PowerProtect DD
- Mounting of a backup as a file system
- Support for sparse-file backups

PowerProtect Data Manager uses traditional file-based backup technology when backing up a specific set of files or directories. During these backups, the entire directory structure of the file system is traversed. These backups take longer to complete than block-based backups.

- NOTE:** Applying an exclusion filter to a protection policy automatically results in a file-based backup. If you are backing up a large file system, it might be more efficient to back up all the data instead. Alternatively, move the assets being filtered to a different protection policy, allowing the remaining unfiltered assets to use a block-based backup.

Understanding backup terminology and managing backup frequency

When scheduling backups in a protection policy, be aware of the following:

- Different protection-policy types can use different terminology to describe available backup levels. This terminology can differ not only between protection-policy types, but also from traditional terminology.
- To avoid high CPU usage that can lead to failure issues, do not schedule backups more often than recommended.

To understand the different backup levels to manage backup frequencies, see the following table.

Table 24. Backup terminology and frequency

Protection-policy types	Available backup levels	Description	Equivalent traditional terminology	Recommended minimum backup interval
VMware application-aware	Full	All the data is backed up.	Full	Monthly
	Synthetic Full	Only the data that has changed since the last synthetic-full or full backup is backup up. An operation to merge these changes with the last synthetic-full or full backup produces a full backup in storage. Only the changed blocks are copied over the	A differential backup is performed, followed by a merge operation that produces a full backup in storage.	12 hours

Table 24. Backup terminology and frequency (continued)

Protection-policy types	Available backup levels	Description	Equivalent traditional terminology	Recommended minimum backup interval
		network, but the result is still a full backup in storage.		
	Log	The transaction logs are backed up.		30 minutes
VMware crash-consistent	Full	All the data is backed up.	Full	Monthly
	Synthetic Full	Only the data that has changed since the last synthetic-full or full backup is backed up. An operation to merge these changes with the last synthetic-full or full backup produces a full backup in storage. Only the changed blocks are copied over the network, but the result is still a full backup in storage.	A differential backup is performed, followed by a merge operation that produces a full backup in storage.	12 hours
Kubernetes crash-consistent	Full	The namespace metadata and persistent volumes are backed up.	Full	Daily
	Synthetic Full	Only the data that has changed for persistent volumes on VMware first-class disks since the last synthetic-full or full backup is backed up. The namespace metadata and all other persistent volumes are backed up in full. Although not all the data is copied over the network, the result is still a full backup in storage.	A combination of full and differential backups are performed, followed by a merge operation that produces a full backup in storage.	12 Hours
File System centralized	Full	All the data is backed up.	Full	Monthly
	Synthetic Full	Only the data that has changed since the last synthetic-full or full backup is backed up. An operation to merge these changes with the last synthetic-full or full backup produces a full backup in storage. Only the changed blocks are copied over the network, but the result is still a full backup in storage.	A differential backup is performed, followed by a merge operation that produces a full backup in storage.	12 hours
Microsoft Exchange Server centralized	Full	All the data is backed up.	Full	Weekly
	Synthetic Full	Only the data that has changed since the last synthetic-full or full backup is backed up. An operation to merge these changes with the last synthetic-full or full backup produces a full backup in storage. Only the changed blocks are copied over the	A differential backup is performed, followed by a merge operation that produces a full backup in storage.	12 hours

Table 24. Backup terminology and frequency (continued)

Protection-policy types	Available backup levels	Description	Equivalent traditional terminology	Recommended minimum backup interval
		network, but the result is still a full backup in storage.		
Microsoft SQL Server centralized	Full	All the data is backed up.	Full	Daily
	Differential	Only the data that has changed since the last differential backup or the last full backup if there are no other differential backups is backed up.	A differential backup is performed, followed by a merge operation that produces a full backup in storage.	12 hours
	Log	The transaction logs are backed up.		30 minutes
Network Attached Storage	Full	All the data is backed up.	Full	Daily ① NOTE: it is recommended to perform a full backup after updating to PowerProtect Data Manager 19.12.
	Synthetic Full	Only the data that has changed since the last synthetic-full or full backup is backup up. An operation to merge these changes with the last synthetic-full or full backup produces a full backup in storage. Only the changed files are copied over the network, but the result is still a full backup in storage.	An incremental backup is performed, followed by a merge operation that produces a full backup in storage.	Daily
Oracle centralized	Full	All the data is backed up.	Full	Daily
	Incremental Cumulative	Only the data that has changed since the last level 0 full backup is backed up.	Differential	12 hours
	Incremental Differential	Only the data that has changed since the last incremental differential backup or the last full backup if there are no other incremental differential backups is backed up.	Incremental	6 hours
	Log	The archived logs are backed up.		30 minutes
SAP HANA centralized	Full	All the data is backed up.	Full	Daily
	Differential	Only the data that has changed since the last full backup is backed up.	Differential	12 hours
	Incremental	Only the data that has changed since the last data backup. The last data backup could be an incremental, differential, or full backup.	Incremental	6 hours

Table 24. Backup terminology and frequency (continued)

Protection-policy types	Available backup levels	Description	Equivalent traditional terminology	Recommended minimum backup interval
VMAX storage group centralized ① NOTE: Not applicable to the PowerProtect Data Manager appliance.	Synthetic Full	Only the data that has changed since the last synthetic-full or full backup is backed up. An operation to merge these changes with the last synthetic-full or full backup produces a full backup in storage. Only the changed blocks are copied over the network, but the result is still a full backup in storage.	A differential backup is performed, followed by a merge operation that produces a full backup in storage.	12 hours
Block Volume	Synthetic Full	Backs up only the blocks that have changed since the last synthetic-full or full backup, and then performs an operation to merge those changes with the last synthetic-full or full backup in order to produce a full backup in storage. Only the changed blocks are copied over the network, but the result is still a full backup in storage.	An incremental backup is performed, followed by a merge operation that produces a full backup in storage.	6 hours
	Snapshot	Saves the state of the volume or volume group and all files and data within it at a particular point in time. Snapshots provide a method of recovery for data that has been corrupted or accidentally deleted. You can use snapshots to restore the volume or volume group to a previous state.		15 minutes

① **NOTE:** In some situations, a full backup might be performed even though a synthetic-full backup was scheduled. Possible reasons for a full backup include the following:

- There is no existing full backup.
- The size of a volume has changed.
- There has been a file path change.
- The asset host has been rebooted.

The backup frequency of log, differential, incremental-cumulative, incremental-differential, and incremental backups cannot be greater than the backup frequency of either full or synthetic-full backups. If you attempt to add or edit a protection policy that uses an invalid backup frequency, PowerProtect Data Manager prevents you from saving the protection policy. You can increase the backup frequency of a protection policy by scheduling more full or synthetic-full backups with different retention times to meet your requirements.

Replication triggers

PowerProtect Data Manager orchestrates protection policy replication objectives independently of the primary backup. When you add a replication objective to a policy, select one of the available triggers.

The default replication trigger is a schedule window that you define by setting a recurrence period plus start and end times. Replication occurs during the defined window. For example, every day between 8 p.m. and 12 a.m.

You can also trigger replication immediately after the completion of the associated primary backup, whether scheduled or manual. At the start of the primary backup, PowerProtect Data Manager generates an associated replication job that remains queued until the end of the protection job. If the backup fails or completes with exception, the associated replication job is skipped. Restarting the protection job queues the associated replication job again.

When you create a replication objective, you can specify either scheduled replication or replication after backup completion, which is applicable to both centralized and self-service protection policies.

NOTE: For replication after backup completion, it is recommended that you update the application agents to the latest version.

Depending on the type of backup, the following versions are required to ensure that replication occurs immediately after the backups complete:

- For self-service primary backups, update all application agents to PowerProtect Data Manager version 19.12 or later.
- For centralized primary backups, update all application agents to PowerProtect Data Manager version 19.11 or later.

If you want to replicate only specific backups, perform a manual replication of these backups in advance.

Using a schedule can help you manage network traffic by replicating during off-peak hours. However, for larger backup sets, the primary backup may not finish before the start of the replication schedule, which creates a replication backlog. Replication after backup completion prevents a replication backlog from forming.

To prevent data loss, the replication after backup completion trigger replicates new backups from the primary objective and any outstanding backups that have not yet replicated.

A job status of Completed with Exceptions during replication

After a triggered replication job, you might see a job status message similar to the following:

Completed with Exceptions

```
ABA0017: plc_linux_rac: Backup was successful for the ORACLE_DATABASE asset ORCLPP on the host oracle.test.com but the copy metadata information is currently unavailable.
```

The backup of this asset completed successfully but the copy metadata information has not yet been discovered by PowerProtect Data Manager. If the 'Replicate immediately upon backup completion' option is enabled for this protection policy, the replication job for the copy might appear in 'Unknown' or 'Cancel' state. Once the copy metadata is discovered by PowerProtect Data Manager, the copy will be replicated.

Review the backup copy details in the View Copies pane of the PowerProtect Data Manager UI Infrastructure > Assets window to determine when the discovery is complete.

If you see this message, the replication backup is not immediately available.

To correct this issue, either wait for the next automatic discovery or initiate a discovery.

Adding or editing a protection policy

You can use the PowerProtect Data Manager user interface to add a protection policy to protect an asset. You can also change the details of an existing protection policy.

Adding a protection policy

You can add a protection policy to protect any of the following asset types. For more information, see the appropriate publication.

Table 25. Protection-policy asset types

Asset type	Publication
File System data	<i>PowerProtect Data Manager File System User Guide</i>
Kubernetes cluster namespaces and PVCs	<i>PowerProtect Data Manager Kubernetes User Guide</i>

Table 25. Protection-policy asset types (continued)

Asset type	Publication
Microsoft Exchange Server databases	<i>PowerProtect Data Manager Microsoft Exchange Server User Guide</i>
Microsoft SQL Server databases	<i>PowerProtect Data Manager Microsoft SQL Server User Guide</i>
Network Attached Storage (NAS) share and appliance data	<i>PowerProtect Data Manager Network-Attached Storage User Guide</i>
Oracle RMAN databases	<i>PowerProtect Data Manager Oracle RMAN User Guide</i>
SAP HANA databases	<i>PowerProtect Data Manager SAP HANA User Guide</i>
Storage Direct data	<i>PowerProtect Data Manager Storage Direct User Guide</i>
Virtual machines	<i>PowerProtect Data Manager Virtual Machine User Guide</i>
Block volumes	<i>PowerProtect Data Manager Storage Array User Guide</i>

Editing a protection policy

You can change any of the following information for an existing enabled or disabled protection policy:

- Policy name and description
- Adding or removing assets from the policy
- Backup and replication schedule
 - ⓘ **NOTE:** You can delete any full or synthetic-full backup schedule except for the first one that was created. The first backup schedule created cannot be deleted.
- Backup optimization mode
- Settings for network interface, storage target, storage unit, retention lock and Service Level Agreement (SLA).

You cannot modify a protection policy type or purpose. For these actions, add a policy. Storage quotas cannot be changed by editing a policy.

- ⓘ **NOTE:** Once you save changes for an enabled or disabled policy, most changes take effect immediately. For a disabled policy's primary backup schedules, however, the changes do not take effect until you reenables the policy, since these schedules do not run in **Disabled** state.

Modify a policy name and description, objectives, or options

The following procedure describes how to change an existing policy name and description, schedule and objectives, or additional backup options in the PowerProtect Data Manager UI.

Prerequisites

If applicable, complete all of the virtual network configuration tasks before you assign any virtual networks to the protection policy.

About this task

- ⓘ **NOTE:** You can also edit a protection policy to add or remove assets. Detailed instructions for adding assets to a policy or removing assets from a policy are provided in the section [Add or remove assets in a protection policy](#).

Steps

1. From the left navigation pane, select **Protection > Protection Policies**.
The **Protection Policies** window appears.
2. Select the protection policy that you want to modify, and click **Edit**.

The **Edit Policy** window opens on the **Summary** page. From this page, you can click edit next to any available row to change specific policy details.

3. In the **Name** or **Description** rows, click **Edit**.

The **Type** page displays.

NOTE: You cannot change the type or purpose of an existing policy.

4. In the **Objectives** row, click **Edit**.

The **Objectives** page displays. From this page, you can change the backup schedule, modify the settings for the network interface, and enable or disable the retention lock.

You can also change the storage targets by selecting a new **Storage Name** in the **Primary Backup** and **Replicate** rows. For more information about changing storage targets, see the section *Changing storage targets*.

5. In the **Options** row, click **Edit**.

The **Options** page displays. From this page, you can change the backup optimization mode (for example, from Performance to Capacity), select whether to include or exclude swap files from the backup, and select whether to quiesce the guest file system during the backup.

NOTE: For virtual machine protection policies, two types of protection mechanisms are used—Transparent Snapshot Data Mover (TSDM), and VMware vStorage API for Data Protection (VADP). Updates to the policy options can result in changes to the protection mechanism used to move virtual machine data. When the protection mechanism changes, a new, full backup is performed, which might take awhile to complete.

6. After making your changes, click **Next** to save the changes and return to the **Summary** page.

7. On the **Summary** page, click **Finish**.

An informational dialog displays.

8. Click **OK** to exit the dialog, or click **Go to Jobs** to open the **Jobs** window to monitor the backup of the new protection policy.

Changing storage targets

A storage target consists of a protection storage system and associated storage unit. You can change the selected storage target elements for each protection policy.

When you edit the primary backup and replication objectives for protection policies:

- The **Storage Name** drop-down list shows the current protection storage system. The drop-down list also contains other protection storage systems that are available. Select **Add** to configure more protection storage.
- The **Storage Unit** drop-down list shows the storage unit that PowerProtect Data Manager targets on the selected protection storage system. From this drop-down list, you can select other storage units under the control of PowerProtect Data Manager. Select **New** to create a storage unit.

When you change the storage target, appropriately configure any dependencies. For example, configure a cloud provider for the updated storage target in the dependent protection policy objective.

NOTE: Network interfaces that exist on a DD 7.4.x or earlier system and that are configured to use an uncompressed IPv6 format cannot be discovered. An example of an uncompressed IPv6 format is 2620:0000:0170:0597:0000:0000:0001:001a. An example of a compressed IPv6 format is 2620:0:170:597::1:1a. To use these network interfaces, reconfigure them to use either an IPv4 address or a compressed IPv6 address, and then initiate a discovery.

Impacts

Changing the primary objective storage target for some asset types may cause skipped backups until the next scheduled full backup:

- VMware virtual machine application-aware
- SAP HANA
- Oracle RMAN

Perform a manual full backup for these policies. Manual backups of protected assets provides instructions.

The following asset types do not require additional action:

- VMware virtual machine crash-consistent

- Kubernetes
- Network Attached Storage (NAS)
- Storage Group
- Microsoft Exchange Server
- Microsoft SQL Server
- File systems

For these asset types, the next backup automatically becomes a full backup.

Replication objectives do not require additional action.

Protection storage

Managing Storage provides more information about working with protection storage, including configuring additional protection storage systems and changing quota settings.

When reviewing the list of selected and available protection storage systems, consider the following:

- It is not recommended that policy objectives share protection storage systems because this configuration does not increase data availability. However, some environments may require replicas with different retention periods, where multiple objectives share a protection storage system.
- Only protection storage that has been licensed and configured for use by the current protection policy appears in the drop-down list.
- Changing protection storage systems for Storage Group protection policies is not supported.

Storage units

Storage units provides more information about working with storage units, including applicable limitations and maintenance considerations.

If you select **New**, PowerProtect Data Manager creates a storage unit for this protection policy. The new storage unit name is based on the protection policy name plus an identifier. Storage units provides more instructions for changing the quota configuration.

You can also select an existing storage unit under the control of PowerProtect Data Manager. The drop-down list displays the available storage units on the selected protection storage system. If the storage unit name is truncated due to space limitations, hover over the list entry to see the full storage unit name and quota information.

Changing storage units for Storage Group protection policies is not supported.

Redeploying storage targets

Redeploying the storage target of a protection policy results in duplicate entries after discovery unless a certain procedure is followed.

To prevent duplicate entries of a redeployed storage target in the **Replication Targets** window, remove it from PowerProtect Data Manager and any relevant protection policies before redeploying it. After the storage target has been redeployed, wait for it to be discovered again, and then add it back to the relevant protection policies.

Replication to shared protection storage

To improve flexibility for external workflows and reduce infrastructure costs, PowerProtect Data Manager supports sharing protection storage across multiple objectives.

To service workflows outside of PowerProtect Data Manager, you may require different retention periods for different replicas. Since retention periods are set at the objective level, configuring different retention periods requires additional replication objectives.

Under most circumstances, additional replication objectives target storage units that reside on different protection storage systems. Replicating to separate protection storage provides additional data availability.

To support external workflows without requiring separate protection storage systems for each additional objective, PowerProtect Data Manager supports targeting different storage units on the same protection storage system. To further

reduce costs, you can target the same protection storage system where the primary backup resides. In this case, the external workflow provides the additional data safety.

NOTE:

Because PowerProtect Data Manager is unaware of external workflows, the UI issues a warning when you configure a policy with multiple objectives that share the same protection storage system. This configuration is uncommon, so verify the storage targets and the use case before you continue.

The UI also issues a warning where the selected storage unit is the source for any MTree replication workflow. This workflow may belong to another application. Verify the storage targets before you continue. These notifications require DDOS 7.7 or later.

Add or remove assets in a protection policy

Perform the following steps in the PowerProtect Data Manager UI to add or remove an asset in a protection policy.

About this task

When a protection policy is edited and new assets are added, backups for the new assets start from the next scheduled FULL backup job for the protection policy.

Steps

1. From the left navigation pane, select **Protection > Protection Policies**.

The **Protection Policies** window appears.

2. Select the protection policy that you want to modify, and click **Edit**.

The **Edit Policy** window opens on the **Summary** page.

3. In the **Assets** row, click **Edit**.

The **Assets** page appears.

NOTE: For virtual machine protection policies, the view that you selected when creating the policy is retained in this page, and cannot be changed. For example, if you set up this policy with **View Asset Table** selected, all assets protected by this policy will display in a table on this page, and the option to select **View by Host** will be disabled. Both views provide additional information about the virtual machines, such as any currently associated tags, protection rules, and whether the virtual machine is already assigned to another policy, to help you identify which assets you want to add or remove from this policy.

4. To remove containers or assets from the protection policy, select the object and click **Remove**.

The **Assets** page updates with the changes.

NOTE: When an asset is moved out of policy while a backup is in progress, PowerProtect Data Manager sets the default retention period for that asset as 30 days. You can modify the retention period for that asset according to your requirement.

5. To add a container or asset to the protection policy:

- a. Click **+ Add**.

The **Add Unprotected Assets** dialog displays any objects that are unprotected.

- b. Select the individual unprotected assets that you want to add to the policy, or select a container level within the hierarchy to add all assets within that level, and then click **Add**.

The **Assets** page updates with the changes.

6. Optionally, if you want to exclude non-production VMDKs such as network shares or test disks from a protection policy:

- a. Select the virtual machine asset from the list, and then click **Manage Exclusions** in the **Disk Excluded** column.

The **Exclude Disks** dialog box appears. By default, the slider next to each VMDK is set to **Included**.

- b. For each disk that you want to exclude, move the slider to the right. The status updates to **Excluded**.

- c. Click **Save**. The **Assets** page updates to indicate the number of disks for that particular asset that will be excluded from the protection policy.

7. Click **Next** to save the changes and go to the **Summary** page.

8. In the **Summary** page, click **Finish**

An informational dialog box appears.

- Click **OK** to exit the dialog box, or click **Go to Jobs** to open the **Jobs** window to monitor the backup of the new protection policy.

Edit the retention period for backup copies


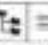
You can edit the retention period of one or more backup copies to extend or shorten the amount of time that backups are retained.

About this task

You can edit retention for all asset types and backup types, except block volumes.

Steps

- From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
- On the **Assets** window, select the tab for the asset type for which you want to edit retention. If a policy has been assigned, the table lists the assets that have been discovered, along with the associated protection policy.

Use the   icons to switch between a tree or folder view of the vCenter server hierarchy, or a list view of all virtual machine assets discovered within the vCenter server.

NOTE: For virtual machine assets, you can click the link in the **Disk Excluded** column next to a virtual machine asset to view VMDKs that have been excluded from the protection policy. You cannot, however, edit disk inclusion or exclusion from this window. To change the disks that are excluded for a protected asset, select the policy from the **Protection Policies** window and click **Edit**.

- Select a protected asset from the table, and then click **View Copies**. The **Copy Locations** pane identifies where the backups are stored.
- In the left pane, click  to the right of the icon for the asset. The table in the right pane lists the backup copies.
- Select one or more backup copies from the table and click **Edit Retention**.

NOTE: If the deletion of an asset backup copy fails, the **Copy Status** is changed from **Available** to another state, and as a result the **Edit Retention** button is disabled. The **Edit Retention** button is enabled only when the **Copy Status** is **Available**.

- Choose one of the following options:
 - To select a calendar date as the expiration date for backups, select **Retention Date**.
 - To define a fixed retention period in days, weeks, months, or years after the backup is performed, select **Retention Value**. For example, you could specify that backups expire after 6 months.

NOTE: When you edit the retention period for copies that are retention locked, you can only extend the retention period.

- When satisfied with the changes, click **Save**.
The asset is displayed in the list with the changes. The **Retention** column displays both the original and new retention period, and indicates whether the retention period has been extended or shortened.

Viewing a summary of protection policies

You can use the PowerProtect Data Manager UI to view a summary of information about a protection policy.

From the left navigation pane, select **Protection > Protection Policies** to view the **Protection Policies** window.

The **Protection Policies** window displays the following columns of information for each protection policy.

- Name**
- Category**
- Asset Type**
- Protected Asset Size**
- Last run status**
- Violations**

• State

The entries in the **Name** and **Last Run Status** columns are links to additional information about the related protection policy.

View assets assigned to a protection policy

You can view assets that are assigned to a protection policy. If assets move from one protection policy to another, you can verify the results from the details window for the protection policy.

About this task

To view the assets that are assigned to a protection policy, complete the following steps:

Steps

1. From the left navigation pane, select **Protection > Protection Policies**.
The **Protection Policies** window opens.
2. Click the name link of the protection policy to view its details.
The details window for the selected protection policy opens and displays information about the policy.
3. Click the asset count link next to **Assets**.
The **Assets** window appears and displays the assets that are assigned to the protection policy.
4. To export asset records for the protection policy, in the **Assets** window, click **Export All**.

View the status of the last-run job of a protection policy

You can use the **Protection Policies** window to determine if the last-run job of a protection policy was successful.

About this task

To view the status of the last-run job of a protection policy, complete the following steps:

Steps

1. From the left navigation pane, select **Protection > Protection Policies**.
The **Protection Policies** window opens.
2. Review the information displayed in the **Last Run Status** column for the protection policy.
3. Optionally, click the last-run status link of the protection policy to view the **Protection Jobs** window for more information about the job.

NOTE: The **Protection Jobs** window displays only the most recently run protection jobs. To view the most recently run system jobs, select **Jobs > System Jobs** from the left navigation pane to view the **System Jobs** window.

Run an asset-protection report

This option enables you to run an asset-protection report and save the report in CSV format so that you can download an Excel file of protection results data.

Steps

1. From the PowerProtect Data Manager UI, select **Protection > Protection Policies**.
2. Select the protection policy for which you would like to export the protection records.
If you do not select a protection policy, PowerProtect Data Manager exports the protection records for all the protection policies.
3. Click **Run Asset Protection Report**.
The **Export Asset Protection** window appears.
4. Specify the following fields for the export:

a. The **Time Range**.

The default is **Last 24 hours**.

This refers to the last complete midnight-to-midnight 24-hour period; that is, yesterday. So, any events that have occurred since the most recent midnight are not in the CSV export. For example, if you run the CSV export at 9am, any events that have occurred in the last 9 hours are not in the CSV export. This is to prevent the overlapping of or partial exporting when queried mid-day on a regular or irregular basis.

b. The **Job Status**.

c. Click **Download**.

If applicable, the navigation window appears for you to select the location to save the .csv file.

5. If applicable, save the .csv file in the desired location and then click **Save**.

Add a service-level agreement

SLA Compliance in the PowerProtect Data Manager UI enables you to add a service-level agreement (SLA) that identifies your service-level objectives (SLOs). You use the SLOs to verify that your protected assets are meeting the service-level agreements (SLAs).

About this task

NOTE: When you create an SLA for Cloud Tier, you can include only full backups in the SLA.

NOTE: The **Extended Retention** SLA only applies to protection policies created in PowerProtect Data Manager 19.11 or earlier. The Extended Retention objective was removed in PowerProtect Data Manager 19.12. Protection policies that were created in earlier releases with the **Extended Retention** SLA are supported. However, you cannot edit the **Extended Retention** SLA in these policies.

In the **SLA Compliance** window, you can export compliance data by using the **Export All** functionality.

Steps

1. From the PowerProtect Data Manager UI, select **Protection > SLA Compliance**.

The **SLA Compliance** window appears.

2. Click **Add** or, if the assets that you want to apply the SLA to are listed, select these assets and then click **Add**.

The **Add Service Level Agreement** wizard appears.

3. Select the type of SLA that you want to add, and then click **Next**.

- **Policy.** If you choose this type, go to step 4.
- **Backup.** If you choose this type, go to step 5.
- **Replication.** If you choose this type, go to step 6.
- **Cloud Tier.** If you choose this type, go to step 7.

You can select only one type of Service Level Agreement.

4. If you selected **Policy**, specify the following fields regarding the purpose of the new Policy SLA:

a. The **SLA Name**.

b. If applicable, select **Minimum Copies**, and specify the number of Backup, Replication, and Cloud Tier copies.

c. If applicable, select **Maximum Copies**, and specify the number of Backup, Replication, and Cloud Tier copies.

d. If applicable, select **Available Location** and select the applicable locations. To add a location, click **Add Location**.

Options include the following:

- **In**—Include locations of all copies in the SLO locations. Selecting this option does not require every SLO location to have a copy.
- **Must In**—Include locations of all copies in the SLO locations. Selecting this option requires every SLO location to have at least one copy.
- **Exclude**—Locations of all copies must be non-SLO locations.

NOTE: Policy files backed up on a storage unit with indefinite retention hold (IRH) enabled cannot be deleted or modified, even after retention lock expiry. It is therefore recommended that you do not select the **Maximum Copies** option because this setting conflicts with IRH. Otherwise, the SLA will not complete successfully once the number of copies exceeds the specified number.

- e. If applicable, select **Allowed in Cloud through Cloud Tier/Cloud DR**.
 - f. Click **Finish**, and then go to step 9.
5. If you selected **Backup**, specify the following fields regarding the purpose of the new **Backup** SLA:
- a. The **SLA Name**.
 - b. If applicable, select **Recovery Point Objective required (RPO)**, and then set the duration. The purpose of an RPO is business continuity planning. It indicates the maximum targeted period in which data (transactions) might be lost from an IT service due to a major incident.

NOTE: You can select only **Recovery Point Objective required** to configure as an independent objective in the SLA, or select both **Recovery Point Objective required** and **Compliance Window for copy type**. If you select both, the RPO setting must be one of the following:

 - Greater than 24 hours or more than the Compliance window duration, in which case RPO validation occurs independent of the Compliance Window.
 - Less than or equal to the Compliance Window duration, in which case RPO validation occurs within the Compliance Window.
 - c. If applicable, select **Compliance Window for copy type**, and then select a schedule level from the list, for example, **All, Full, Cumulative**, and set the duration. **Duration** indicates the amount of time necessary to create the backup copy. Ensure that the **Start Time** and **End Time** of backup copy creation falls within the Compliance Window duration specified.

This window specifies the time during which you expect the specified activity to take place. Any specified activity that occurs outside of this **Start Time** and **End Time** triggers an alert.
 - d. If applicable, select the **Verify expired copies are deleted** option.

Verify expired copies are deleted is a compliance check to see if PowerProtect Data Manager is deleting expired copies. This option is disabled by default.

NOTE: Data that is backed up on a storage unit with IRH enabled cannot be deleted or modified, even after retention lock expiry. It is therefore recommended that you do not select the **Verify expired copies are deleted** option because this setting conflicts with IRH. Otherwise, the SLA will not complete successfully.
 - e. If applicable, select **Retention Time Objective**, and specify the number of Days, Months, Weeks, or Years.

NOTE: The value of **Retention Time Objective** must match the lowest retention value of the backup levels of the target objectives of this policy. For example, if the synthetic full backup **Retain For** is 30 days but the full backup **Retain For** is 60 days, set the **Retention Time Objective** to 30 days.
 - f. If applicable, select the **Verify Retention Lock is enabled for all copies** option. This option is disabled by default.
 - g. Click **Finish**, and go to step 9.

The **SLA Compliance** window appears with the new SLA.
6. If you selected **Replication**, specify the following fields regarding the purpose of the new Replication SLA:
- a. The **SLA Name**.
 - b. If applicable, select the **Compliance Window**, and specify the **Start Time** and **End Time**.



This window specifies the times that are permissible and during which you can expect the specified activity to occur. Any specified activity that occurs outside of this start time and end time triggers an alert.
 - c. If applicable, select the **Verify expired copies are deleted** option.

Verify expired copies are deleted is a compliance check to see if PowerProtect Data Manager is deleting expired copies. This option is disabled by default.


NOTE: Data that is replicated on a storage unit with IRH enabled cannot be deleted or modified, even after retention lock expiry. It is therefore recommended that you do not select the **Verify expired copies are deleted** option because this setting conflicts with IRH. Otherwise, the SLA will not complete successfully.
 - d. If applicable, select **Retention Time Objective**, and specify the number of Days, Months, Weeks, or Years.

NOTE: Set the value of **Retention Time Objective** to match the lowest retention value of the backup levels of the target objectives of this policy.
 - e. If applicable, select the **Verify Retention Lock is enabled for all copies** option. This option is disabled by default.
 - f. Click **Finish**, and go to step 9.

The **SLA Compliance** window appears with the newly added SLA.
7. If you selected **Cloud Tier** type SLA, specify the following fields regarding the purpose of the new Cloud Tier SLA:
- a. The **SLA Name**.

- b. If applicable, select the **Verify expired copies are deleted** option.
This option is a compliance check to determine if PowerProtect Data Manager is deleting expired copies. This option is disabled by default.
 - c. If applicable, select **Retention Time Objective** and specify the number of Days, Months, Weeks, or Years.
 **NOTE:** Set the value of **Retention Time Objective** to match the lowest retention value of the backup levels of the target objectives of this policy.
 - d. If applicable, select the **Verify Retention Lock is enabled for all copies** option. This option is disabled by default.
 - e. Click **Finish**.
8. If the SLA has not already been applied to a protection policy:
 - a. Go to **Protection > Protection Policies**.
 - b. Select the policy, and then click **Edit**.
 9. In the **Objectives** row of the **Summary** window, click **Edit**.
 10. Do one of the following, and then click **Next**:
 - Select the added Policy SLA from the **Set Policy Level SLA** list.
 - Create and add the SLA policy from the **Set Policy Level SLA** list.
 The **Summary** window appears.
 11. Click **Finish**.
An informational message appears to confirm that PowerProtect Data Manager has saved the protection policy.
 12. Click **Go to Jobs** to open the **Jobs** window to monitor the backup and compliance results, or click **OK** to exit.
 **NOTE:** Compliance checks occur automatically every day at 2 a.m. Coordinated Universal Time (UTC). If any objectives are out of compliance, an alert is generated at 2 a.m. UTC. The **Validate** job in the **System Jobs** window indicates the results of the daily compliance check.

For a backup SLA with a required RPO setting that is less than 24 hours, PowerProtect Data Manager performs real-time compliance checks. If you selected **Compliance Window for copy type** and set the backup level to **All**, the real-time compliance check occurs every 15 minutes only within the compliance window. If the backup level is not **All**, or if a compliance window is not specified, the real-time compliance check occurs every 15 minutes without stop.

 **NOTE:** If the backup SLA has a required RPO setting of 24 hours or greater, compliance checks occur daily at 2 a.m. UTC. Real-time compliance checks do not occur for backup SLAs with an RPO setting of 24 hours or greater.

Real-time compliance-check behavior

If the asset was not backed up within the RPO backup interval requirement, an alert indicates that the RPO of the asset is out of compliance. This alert is generated once within an RPO period. If the same backup copy is missed when the next compliance check occurs, no further alerts are generated.

If the asset was backed up within the RPO backup interval requirement, the RPO of the asset is in compliance.

If multiple assets in a policy are out of compliance, a single alert is generated. This alert includes information for all the assets in the policy. In the **Alerts** window, the asset count next to the alert summary indicates the number of assets that are out of compliance in the policy.

13. In the **Jobs** window, click  next to an entry to view details on the SLA Compliance result.

Run a compliance report

This option enables you to run a compliance report and save the report in CSV format so that you can download an Excel file of compliance results data.

Steps

1. From the PowerProtect Data Manager UI, select **Protection > SLA Compliance**.
The **SLA Compliance** window appears. The PowerProtect Data Manager **SLA Compliance** window displays the following information:
 - SLA Name
 - Stage Type
 - Policies At Risk
 - Objectives Out of Compliance

- Impacted Assets
2. Select the SLA for which you would like to export the compliance records.
 3. Click **Run Compliance Report**.
The **Run Compliance Report** window appears.
 4. Specify the following fields for the export:
 - a. The **Time Range**.
The default is **Last 24 hours**.
This refers to the last complete midnight-to-midnight 24 hour period; that is, yesterday. So, any events that have occurred since the most recent midnight are not included in the CSV export. For example, if you run the CSV export at 9am, any events that have occurred in the last 9 hours are not included in the CSV export. This is to prevent the overlapping of or partial exporting when queried mid-day on a regular or irregular basis.
 - b. The **Job Status**.
 - c. Click **Download .CSV**.
If applicable, the navigation window appears for you to select the location to save the .csv file.
 5. If applicable, save the .csv file in the desired location and click **Save**.

Disable a protection policy

From the PowerProtect Data Manager UI, you can disable a protection policy to temporarily stop running certain backup objectives of this policy.

About this task

There are several reasons why you might want to disable a protection policy. For example, by disabling a policy, you can:

- Edit the policy and determine the impact of your changes before these changes take effect.
- Stop backup activity on primary storage if the storage is in maintenance or is temporarily unavailable (for example, during a storage upgrade).

By default, disabling a centralized protection policy stops the primary backup objectives of this policy, including synthetic full backups, full backups, and so on. Any replication and cloud tier objectives, however, continue to run while the policy is disabled. You can also perform manual primary backups of a policy that is in **Disabled** state by using the **Protect Now** functionality in the PowerProtect Data Manager UI.

You can modify the default behavior to make changes regarding which jobs continue to run when a policy is disabled by using System Level overwrites in the REST API. The PowerProtect Data Manager Public REST API documentation provides instructions.

When a protection policy is disabled, you can edit the policy in the same manner that you would edit an enabled policy. The advantage of editing a policy in **Disabled** state is that you can preview the changes before resuming primary backups of the policy. Adding or editing a protection policy provides more information about modifying the details of an existing policy.

Steps

1. From the left navigation pane, select **Protection > Protection Policies**.
The **Protection Policies** window opens.
2. Select one or more policies in **Enabled** state. You can also select the checkbox at the top of the table to select all policies on the current page.
3. Click **Disable**.

Results

The policy status changes to **Disabled**. In **Disabled** state:

- In progress primary backup jobs that are associated with this policy continue to run until complete. If primary backups are scheduled to run during the time that the policy is disabled, those backups do not run, even when you enable the policy again. When you re-enable the policy, future scheduled backups resume.
- All other protection jobs for the policy continue to run according to schedule, unless no primary backup copy exists for the policy. In this case, protection jobs are skipped.
- Manual backups of primary objectives can still be performed.

Protection jobs running for a disabled policy

When a protection policy is disabled, only protection jobs related to the primary backup objectives stop running.

The following table provides information about the types of protection jobs that continue to run when a policy is in **Disabled** state. The column **System level overwrite?** indicates whether the default behavior for this job can be overwritten by using the API command. Note, however, that when a policy is disabled, the setting for at least one of these jobs must remain disabled.

NOTE: If no primary backup copy exists for the disabled policy, other scheduled protection jobs such as replication will display as **Skipped** in the **Protection Jobs** window of the PowerProtect Data Manager UI.

Table 26. Protection jobs running when a policy is disabled

Job category	Purpose	Runs when policy is disabled?	System level overwrite?
Centralized scheduled primary protection	Create a primary backup	No	Yes
Manual backup and replication (Protect Now, Replicate Now)	<ul style="list-style-type: none">Create a primary backup (Protect Now)Replicates primary backup (Replicate Now)	Yes	No
Self-service protection	Create a primary backup	Yes	No
Policy and asset configuration	Prepare for protection or copy management jobs	Yes	No
Replication	Copy management (location)	Yes	Yes
Cloud DR	Copy management (location)	Yes	Yes
Extended Retention	Copy management (retention)	Yes	Yes
Cloud Tier	Copy management (location)	Yes	Yes
SLA compliance verification	Copy management (report and alert)	Yes	Yes
Delete expired copy	Copy management (reclaiming space on DD)	Yes	Yes

Enable a disabled protection policy

To reenable a disabled policy, perform the following steps:

Steps

- From the PowerProtect Data Manager UI, select **Protection > Protection Policies**.
- Select one or more policies in **Disabled** state. You can also select the checkbox at the top of the table to select all policies on the current page.
- Click **Enable**.

Results

The status changes to **Enabled**. Primary backups for the reenabled policies resume according to the protection policy schedule.

Customize the default behavior of disabled policies

By default, a protection policy in **Disabled** state prevents the primary backup objectives of this policy from running, but does not stop other protection jobs. You can, however, change the default behavior to also stop other activities, such as replication and cloud tiering, by using the REST API.

The PowerProtect Data Manager Public REST API documentation provides instructions.

Delete a protection policy

Perform the following steps to delete a protection policy that is not protecting any assets.

Prerequisites

If the policy you want to delete protects assets, you must associate those assets with a different protection policy before you can delete the policy.

Steps

1. From the PowerProtect Data Manager UI, select **Protection > Protection Policies**.
2. Select the policy that you want to delete, and then click **Delete**.

Results

After you delete a policy, clean-up of unnecessary components within protection storage occurs automatically according to schedule. Clean-up includes storage units under the control of PowerProtect Data Manager and the corresponding DD Boost users, according to the rules for storage units.

Overview of PowerProtect Data Manager Cloud Tier

The PowerProtect Data Manager Cloud Tier feature works in tandem with the Cloud Tier feature of DD systems to move PowerProtect Data Manager backups to the cloud. This provides long-term storage of PowerProtect Data Manager backups by seamlessly and securely tiering data to the cloud.

From the PowerProtect Data Manager UI, you configure Cloud Tier to move PowerProtect Data Manager backups from protection storage to the cloud, and you can perform seamless recovery of these backups.

Cloud storage units must be pre-configured on the protection storage system before they are configured for Cloud Tier in the PowerProtect Data Manager UI. The *DDOS Administration Guide* provides further information.

Add a Cloud Tier objective to a protection policy

For some protection policy types, you can add a Cloud Tier objective to a protection policy in order to move local full backups to Cloud Tier after a predefined number of days.

Prerequisites

- Ensure that a protection storage system is set up for Cloud Tiering, with the system passphrase set.
- Cloud storage units must be pre-configured on the protection storage system before they are configured for Cloud Tier in the PowerProtect Data Manager UI.
- A data movement schedule must be configured and running on the cloud storage unit.
- The Cloud Tier objective can be added to the **Primary Backup** and **Replicate** objectives. The **Primary Backup** and **Replicate** objectives should be using the protection storage system that is set up for Cloud Tiering.

About this task

Cloud Tiering happens at 00:00 UTC each day. Depending on your time zone, this time may be within business hours and thus Cloud Tiering may impact available network bandwidth. Cloud Tiering applies to both centralized and self-service protection policies.

Steps

1. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.
2. From the PowerProtect Data Manager UI, select **Protection > Protection Policies**, and then click **Add**. The **Add Policy** wizard appears.
3. On the **Type** page, enter a name and description, select the type of system to back up, and click **Next**.

The following protection policy types support Cloud Tiering:

- Virtual machine
 - Microsoft SQL Server
 - Microsoft Exchange Server
 - Network Attached Storage (NAS)
 - Oracle
 - SAP HANA
 - File System
 - Kubernetes
 - Block volume
4. On the **Purpose** page, select from the available options to indicate the purpose of the new protection policy, and then click **Next**.
 5. On the **Assets** page, select the assets that you want to protect with this policy, and then click **Next**.
 6. On the **Objectives** page, click **Add** under **Primary Backup** if the primary backup objective is not already created, and fill out the fields in the Target and Schedules panes on the **Add Primary Backup** dialog.

NOTE: There is no minimum recurrence required for the Cloud objective, however, the Cloud Tier objective requires a minimum retention period of 14 days in the **Retain for** field.
 7. Click **Cloud Tier** next to **Primary Backup** or, if adding a Cloud objective for a replication objective that you have added, click **Cloud Tier** under **Replicate**.
An entry for **Cloud Tier** is created to the right of the primary backup objective, or below the replication objective.
 8. Under the entry for **Cloud Tier**, click **Add**.
The **Add Cloud Tier Backup** dialog appears, with summary information for the parent node. This information indicates whether you are adding this Cloud Tier objective for the primary backup objective or the replication objective.
 9. In the **Add Cloud Tier Backup** dialog box, set the following parameters and then click **Save**:
 - Select one or more of the upstream full backups.
 - Select the appropriate Cloud Unit from the **Cloud Target** list.
 - For **Tier After**, set a time of 14 days or more.

The protection policy is now enabled with Cloud Tiering.

NOTE: If the retention period of a copy is less than the time specified in the **Tier After** field, and you do not edit the **Retain for** value of this schedule or its copy to a value greater than the **Tier After** field before the retention period of the copy expires, the copy will not be cloud tiered.
 10. Click **Next** to proceed with the remaining pages of the **Add Policy** wizard, verify the information, and then click **Finish**.
A new job is created, which you can view under the **Jobs** tab after the job completes.

Manage Cloud Tier asset copies

You can manage Cloud Tier copies of assets by changing copy retention time, deleting copies, and recalling copies.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
2. Select an asset and click **View Copies**.
3. Click an asset copy icon.
Cloud Tier backups are listed by cloud storage in the **Location** column.
4. To change how long copies remain in cloud storage, complete the following steps:
 - a. Select a Cloud Tier backup and click **Edit Retention**.
 - b. Choose one of the following options:
 - To select a calendar date as the expiration date for backups, select **Retention Date**.
 - To define a fixed retention period in days, weeks, months, or years after the backup is performed, select **Retention Value**. For example, you could specify that backups expire after 6 months.
 - c. When satisfied with the changes, click **Save**.
The asset is displayed in the list with the changes. The **Retention** column displays both the original and new retention period, and indicates whether the retention period has been extended or shortened.

NOTE: When you edit the retention period for copies that are retention locked, you can only extend the retention period.

5. To delete the copy in cloud storage, select a Cloud Tier backup and click **Delete**. To delete the copy records from the PowerProtect Data Manager database while the copy remains in the protection storage, select **Remove from PowerProtect**.

Delete backup copies and Remove Exchange, File System, Kubernetes, Block Volume, and SQL backup copies from the PowerProtect Data Manager database provides more information.

6. Select a Cloud Tier backup and click **Recall from Cloud** to return the cloud backup to your local protection storage for recovery or backup.

NOTE: If you use Amazon's network to copy data from AWS storage, Amazon charges you for the data transfer.

7. To extend the date to re-tier the copy back to the cloud, select **Edit Recall Retention**.
8. To manually move a copy back to cloud storage, select **Retier**.

Restore Cloud Tier backups to protection storage

Once a Cloud Tier backup is recalled, restore operations of these backups are identical to normal restore operations.

The PowerProtect Data Manager software recalls a copy of the backup from the Cloud Unit to the local (active) tier of protection storage, which then allows you to perform a restore of the backup from the active tier to the client. The status appears as **Cloud**, and changes to **Local Recalled** after cloud recall completes. After the restore, the backup copy is removed from Cloud Tier, and is stored on the active tier of protection storage for a minimum of 14 days, after which the backup may be returned to the cloud depending on your protection policy.

Recall and restore from Cloud Tier

Perform the following steps to recall a backup on Cloud Tier to the active tier on protection storage and restore this backup.

Prerequisites

NOTE: When a backup is recalled from Cloud Tier to the active tier, the copy is removed from Cloud Tier.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
2. On the **Assets** window, select the tab that contains the asset you want to recall from Cloud Tier, and then click **View Copies**.
3. Click **DD**, and then select from one of the available copies that appear in the table.
4. Click **Recall**.
The **Recall from Cloud** dialog box appears.
5. In the **Retain until** box, specify how long you want to keep the copy on the active tier, and then click **OK**.
6. Go to the **Jobs** window to monitor the recall operation.
When the copy has been moved successfully, the **Location** changes from Cloud to Local.
7. Select **Restore > Assets**, and then select the tab that contains the recalled asset.
8. Select the recalled asset, and then click **Restore**.

NOTE: If you are unsure whether the asset has been recalled, click **View Copies** and select **DD** to view the available backup copies. If the asset backup is a recalled copy, the Status column indicates **Local Recalled**.

9. Select the recalled copy to re-tier the copy to the active tier.

Extended retention for protection policies created in PowerProtect Data Manager 19.11 or earlier

NOTE: This section only applies to protection policies created in PowerProtect Data Manager 19.11 or earlier. For protection policies created in PowerProtect Data Manager 19.12 or later, you add multiple full schedules for primary backup and replication objectives. Protection policies that were created in an earlier release with the **Extend Retention** objective are

supported. However, you cannot edit existing extended retention objectives or add new extended retention objectives in these policies. Knowledge Base article 000204454 at <https://www.dell.com/support/> provides detailed information about specific **Extend Retention** objective migration scenarios when updating PowerProtect Data Manager.

For protection policies created in PowerProtect Data Manager 19.11 or earlier, the **Extend Retention** objective allows you to extend the retention period for the primary backup copy for long-term retention. For example, your regular schedule for daily backups uses a retention period of 30 days. However, you can extend the retention period to keep full backups taken on Mondays for 10 weeks.

Both centralized and self-service protection policies support weekly, monthly, and yearly recurrence schedules to meet the demands of your compliance objectives. For example, you can retain the last full backup containing the last transaction of a fiscal year for 10 years. Extended retention periods can retain scheduled full backups with a repeating pattern for a specified amount of time.

For example:

- Retain full yearly backups that are set to repeat on the first day of January for 5 years.
- Retain full monthly backups that are set to repeat on the last day of every month for 1 year.
- Retain full yearly backups that are set to repeat on the third Monday of December for 7 years.

Preferred alternatives

When you define an extended retention objective for a protection policy, you define matching criteria that select preferred backups to retain. If the matching criteria do not identify a matching backup, PowerProtect Data Manager automatically retains the preferred alternative backup according to one of the following methods:

- **Look-back**—Retain the last available full backup that was taken before the matching criteria.
- **Look-forward**—Retain the next available full backup that was taken after the matching criteria.

For example, consider a situation where you configured a protection policy to retain the daily backup for the last day of the month to extended retention. However, a network issue caused that backup to fail. In this case, look-back matching retains the backup that was taken the previous day, while look-forward matching retains the backup that was taken the following day.

By default, PowerProtect Data Manager uses look-back matching to select the preferred alternative backup. A grace period defines how far PowerProtect Data Manager can look in the configured direction for an alternative backup. If PowerProtect Data Manager cannot find an alternative backup within the grace period, extended retention fails.

You can use the REST API to change the matching method or the grace period for look-forward matching. The PowerProtect Data Manager Public REST API documentation provides instructions. If there are no available backups for the defined matching period, you can change the matching method to a different backup.

For look-forward matching, the next available backup can be a manual backup or the next scheduled backup.

Selecting backups by weekday

This section applies to centralized protection policies. Self-service protection policies have no primary backup objective configuration.

When you configure extended retention to match backups by weekday, PowerProtect Data Manager might identify a backup as having been taken on the wrong weekday. This behavior happens where the backup window does not align with the start of the day. PowerProtect Data Manager identifies backups according to the day on which the corresponding backup window started, rather than the start of the backup itself.

For example, consider a backup schedule with an 8:00 p.m. to 6:00 a.m. backup window:

- Backups that start at 12:00 a.m. on Sunday and end at 6:00 a.m. on Sunday are identified as Saturday backups, since the backup window started on Saturday.
- Backups that start at 8:01 p.m. on Sunday and end at 12:00 a.m. on Monday are identified as Sunday backups, since the backup window started on Sunday.
- Backups that start at 12:00 a.m. on Monday and end at 6:00 a.m. on Monday are identified as Sunday backups, since the backup window started on Sunday.

In this example, when you select Sunday backups for extended retention, PowerProtect Data Manager does not retain backups that were taken between 12:00 a.m. and 8:00 p.m. This behavior happens even though the backups occurred on Sunday. Instead, PowerProtect Data Manager selects the first available backup that started after 8:00 p.m. on Sunday for extended retention.

If no backups were created between 8:01 p.m. on Sunday and 6:00 a.m. on Monday, PowerProtect Data Manager retains the next alternative to extended retention. In this example, the alternative was taken after 6:00 a.m. on Monday.

Extended retention backup behavior

When PowerProtect Data Manager identifies a matching backup, automatically extended retention creates a job at the beginning of the backup window for the primary objective. This job remains queued until the end of the backup window.

The following examples describe the behavior of backups with extended retention for centralized and self-service protection.

Centralized protection

An hourly primary backup schedule starts on Sunday at 8:00 p.m., and ends on Monday at 6:00 p.m. with a weekly extended retention objective set to repeat every Sunday. PowerProtect Data Manager selects the first available backup starting after 8:00 p.m. on Sunday for long-term retention.

The following diagram illustrates the behavior of backups with extended retention for a configured protection policy. In this example, full daily backups starting at 10:00 p.m. and ending at 6:00 a.m. are kept for 1 week. Full weekly backups are set to repeat every Sunday and are kept for 1 month.




Figure 2. Extend retention backup behavior

Self-service protection

For self-service backups, PowerProtect Data Manager uses a default backup window of 24 hours. A backup schedule starts on Sunday at 12:00 p.m. and ends on Monday at 12:00 p.m. with a weekly extended retention objective set to repeat every Sunday. PowerProtect Data Manager selects the first available backup that is taken between 12:00 p.m. on Sunday and 12:00 p.m. on Monday for long-term retention.

Replication of extended retention backups

You can change the retention time of selected full primary backups in a replication objective by adding a replication objective to the extended retention backup. The rules in the extended retention objective define the selected full primary backups. Review the following information about replication of extended retention backups.

- Before you configure replication of extended retention backups, create a replication objective for the primary backup.
- Configure the replication objective of the extended retention and match this objective with one of the existing replication objectives based on the primary backup. Any changes to a new or existing storage unit in the extended retention replication objective or the replication objective of the primary backup is applied to both replication objectives.
- The replication objective of extended retention backups only updates the retention time of replicated backup copies. New backup copies are not created in the replication storage.

Manual backups of protected assets

Once assets have been added to a protection policy, you can perform manual backups by using the **Protect Now** functionality in the PowerProtect Data Manager UI.

About this task

You can use a single manual backup from the **Protection > Protection Policies** window to back up multiple assets that are protected in the designated protection policy. The protection policy can be enabled or disabled, but its purpose must not be Exclusion or Self-Service Protection.

When a virtual machine is part of an application-aware protection policy, the manual backup is a full application-aware backup.

The manual backup is managed by other configured objectives (replication, Cloud Tier, Cloud DR) of the parent protection policy. Other properties, such as retention lock, storage target, quotas, and network interfaces, are inherited from the parent protection policy. Jobs managed by this protection policy, such as replication, cloud tiering, and Cloud DR, continue to run after the manual backup job completes.

Steps

1. From the left navigation pane, select **Protection > Protection Policies**.
The **Protection Policies** window appears.
2. Select the protection policy that contains the assets that you want to back up, and click **Protect Now**.
The **Protect Now** wizard appears.
3. On the **Assets Selection** page, select whether you want to back up all assets or choose individual assets that are defined in the protection policy, and then click **Next**.
If you selected the option to choose individual assets for manual backup instead of all assets, the **Assets** page appears with the individual assets available for selection.
 - a. Select the assets that you want to include in the manual backup, and then click **Next** to display the **Configuration** page. If you selected to back up all assets, the **Configuration** page appears.
4. On the **Configuration** page, select **Back up now**, and then select from the available backup types.
5. Edit the retention period if you want to change the default settings, and then click **Next**.
The default settings are inherited from the primary backup objective of the parent protection policy.
6. You can select **Troubleshooting mode** to enable debug logging, and then select the level of logging to use:
 - **Info**—Includes information such as status changes. This is the default log level for scheduled backups and restores.
 - **Debug**—Additional information that helps with problem diagnosis.
 - **Trace**—The most detailed amount of information for complex problem diagnosis.
7. On the **Summary** page, review the settings and then click **Protect Now**.
A notification appears indicating whether the request was processed successfully.

Manual backups of a single protected asset

You can also perform a manual backup from the **Infrastructure > Assets** window, but only for one asset at a time.

About this task

Review the information at Manual backups of protected assets. The protection policy can be enabled or disabled, but its purpose must not be Exclusion or Self-Service Protection. This task creates a full backup for the selected asset.

Steps

1. From the left navigation pane, select **Infrastructure > Assets**.
The **Assets** window appears.
2. Select the tab for the asset type you want to back up.
A list of assets appears.
3. Select an asset from the table that has an associated protection policy.
4. Click **Back Up Now**.
A notification appears indicating whether the request was processed successfully.

Manual replication of protected assets

You can replicate one or more protected assets within a protection policy by using the **Protect Now** functionality in the PowerProtect Data Manager UI. Replication can include all assets which are defined on the protection policy or a subset of these assets. After you select assets, you can replicate all backups or a subset of backups.

Prerequisites

The protection policy purpose must not be **Exclusion**, and the policy must already be configured with a replication objective. You can only manually replicate the replication objectives for the primary backup.



NOTE: VMAX storage groups only support MTree replication, which is performed and scheduled from the DD system. Therefore, manual replication for assets in a VMAX storage group is not supported.

About this task

Replicating a subset of backups is useful when the replication backlog is too large to catch up. For example, when the destination was offline for an extended period or where bandwidth and capacity issues prevent a full replication during the available window.

If the backlog is too large, you can ensure that the destination receives the most recent backups first. You can also reduce the backlog by skipping the future replication of backups that are too old to match the selection criteria.

Steps

1. From the PowerProtect Data Manager UI, select **Protection > Protection Policies**.
2. Select the protection policy that contains the assets that you want to replicate, and click **Protect Now**. The **Protect Now** wizard opens to the **Assets Selection** page.
3. Select whether you want to replicate **All Assets** or a **Custom** selection of assets.
 - If you selected **All Assets**, click **Next**.
 - If you selected **Custom**, a list appears from which you can select individual assets. You can see these assets in tree view or list view.
 - a. Select the assets that you want manually replicate, and then click **Next**. The **Configuration** page appears.
4. Select **Replicate Now**.
5. Select a destination storage target from the **Storage Name** and **Storage Unit** drop-down lists. The selection of storage system and storage unit from these drop-down lists corresponds to the associated replication objective for the primary backup. In some cases, a protection storage system may have more than one storage unit for this policy. The wizard loads the default settings from the protection policy.
6. If you want to change the default settings:
 - You can configure different retention periods for all applicable backup types, or configure the same retention period for all backup types.
 - The default retention period settings are inherited from the settings in the corresponding replication objective of the protection policy.
 - For VMDM, File System, Microsoft Exchange Server, and NAS assets, the retention period for full and synthetic full backups should be the same value.
 - a. Select or clear **Set the same retention time for all replicated copies**.
 - b. Edit the retention period for all applicable backup types.
 - c. Resolve any conflicts or errors, as indicated by the  and  symbols.
7. Select whether you want to replicate **All Copies** or a **Custom** subset of backups. If you selected **Custom**, additional options appear:
 - a. To replicate recent backups by time, select the first option and then type the number of days.
 - b. To replicate a specific number of recent backups, select the second option and then type the number of backups.
 - c. (Optional) To remove all nonmatching backups from the replication backlog for this objective, select **Do not replicate copies outside the selection and mark them as skipped**. PowerProtect Data Manager excludes any skipped backups from future replication activities by this objective. This decision is permanent and the wizard prompts for confirmation.

If the chain for the selected backups has not already replicated, the resulting activity replicates the chain from the last full backup to the selection.

NOTE: Manual replication of a FULL backup for any asset type with a dependency chain (for example, a backup that includes transaction logs) is skipped when the FULL copy has already been replicated, even if the dependencies have not yet been replicated. To replicate any of these dependencies in the backup chain, wait for the scheduled replication, or perform a manual replication with the **All Copies** option selected instead of the **Custom** option.


- (Optional) Click **Select Replication** and then repeat the previous steps to configure manual replication for additional replication policy objectives.
- Click **Next**.
- On the **Summary** page, review the settings and then click **Protect Now**.
A notification appears indicating whether the request was processed successfully.

Manual Cloud Tiering of protected assets

Once you add assets to a protection policy that contains a Cloud Tier objective, you can perform a manual tiering of these assets by using the PowerProtect Data Manager UI.

NOTE: Manual Cloud Tiering of a copy set requires the related protection policy to have a Cloud Tier objective.

To perform on-demand Cloud Tiering:

- From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
- On the **Assets** window, select the tab for the asset type you want to tier. A list of assets appears.
- Select an asset from the table that has an associated protection policy, and then click **View Copies**.
NOTE: You can only select one asset at a time, and the protection policy that is associated with the asset cannot be an exclusion policy.
- In the left pane, click  to the right of the icon for the asset to display the available backup copies in the right pane.
- Select a backup copy, and then click **Tier**. A notification appears indicating whether the request was processed successfully.

Go to the **Jobs** window to monitor the progress of the tiering operation.

Delete backup copies


In addition to deleting backups after the retention period expires, PowerProtect Data Manager enables you to manually delete backup copies from protection storage.

About this task

If you no longer require a backup copy and the retention lock is not enabled, you can delete backup copies prior to their expiration date.

You can perform a backup copy deletion that deletes only a specified part of a backup copy chain, without impacting the ability to restore other backup copies in the chain. When you select a specific backup copy for deletion, only that backup copy and the backup copies that depend on the selected backup copy are deleted. For example, when you select to delete a full backup copy, any other backup copies that depend on the full backup copy are also deleted.

Steps

- From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
- From the **Assets** window, select the tab for the asset type for which you want to delete copies. If a policy has been assigned, the table lists the assets that have been discovered, along with the associated protection policy.
- Select a protected asset from the table, and then click **View Copies**. The **Copy Locations** pane identifies where the backups are stored.
- In the left pane, click  to the right of the icon for the asset. The table in the right pane lists the backup copies.
- Select one or more copies from the table that you want to delete from the DD system, and then click **Delete**.
A preview window opens and displays the selected backup copies.

NOTE: For assets with backup copies that are chained together such as Microsoft SQL Server databases, Oracle databases, SAP HANA databases, and application-aware virtual machines, the preview window lists all the backup copies that depend on the specified backup copy. If you delete a backup copy, PowerProtect Data Manager deletes the specified backup copy and all backup copies that depend on the specified backup copy.

- For all asset types, you can choose to keep the latest backup copies or delete them. By default, PowerProtect Data Manager keeps the latest backup copies. To delete the latest backup copies, clear the checkbox next to **Include latest copies**. For VMAX storage group backup copies, you can choose to delete copies that are grouped together in the same protection transaction or delete only selected copies. By default, PowerProtect Data Manager deletes copies that are grouped together in the same protection transaction. To delete only selected copies, clear the checkbox next to **Include copies in the same protection transaction**.
- To delete the backup copies, in the preview window, click **Delete**.

NOTE: The delete operation may take a few minutes and cannot be undone.

An informational dialog box opens to confirm the copies are being deleted. To monitor the progress of the operation, click **Go to Jobs**. To view the list of backup copies and their status, click **OK**.

NOTE: If the data deletion is successful but the catalog deletion is unsuccessful, then the overall deletion job status appears as **Completed with Exceptions**.

When the job completes, the task summary provides details of each deleted backup copy, including the time that each copy was created, the backup level, and the retention time. The time of copy creation and the retention time are shown in UTC.


An audit log is also generated and provides details of each deleted backup copy, including the time that each copy was created, the backup level, and the retention time. The time of copy creation and the retention time are shown in UTC. Go to **Alerts > Audit Logs** to view the audit log.

- Verify that the copies are deleted successfully from protection storage. If the deletion is successful, the deleted copies no longer appear in the table.

Retry a failed backup copy deletion

If a backup copy is not deleted successfully, you can manually retry the operation.

Steps

- From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
- From the **Assets** window, select the tab for the asset type for which you want to delete copies. If a policy has been assigned, the table lists the assets that have been discovered, along with the associated protection policy.
- Select a protected asset from the table, and then click **View Copies**. The **Copy Locations** pane identifies where the backups are stored.
- In the left pane, click  to the right of the icon for the asset. The table in the right pane lists the backup copies.
- Select one or more backup copies with the **Deletion Failed** status from the table, and then click **Delete**. You can also filter and sort the list of backup copies by status in the **Copy Status** column. The system displays a warning to confirm that you want to delete the selected backup copies.
- Click **OK**. An informational dialog box opens to confirm that the copies are being deleted. To monitor the progress of the operation, click **Go to Jobs**. To view the list of backup copies and their status, click **OK**.
- Verify that the copies are successfully deleted from protection storage. If the deletion is successful, the deleted copies no longer appear in the table.


Export data for deleted Oracle, SAP HANA and Storage Direct backup copies

This option enables you to export results of deleted backup copies to a .csv file so that you can download an Excel file of the data.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
2. From the **Assets** window, select the tab for the asset type for which you want to export results of deleted backup copies. If a policy has been assigned, the table lists the assets that have been discovered, along with the associated protection policy.
3. Select one or more protected assets from the table, and then select **More Actions > Export Deleted Copies**.
If you do not select an asset, PowerProtect Data Manager exports the data for deleted backup copies for all assets for the specific asset type.
4. Specify the following fields for the export:
 - a. **Time Range**
The default is **Last 24 Hours**.
 - b. **Copy Status**
In order to export data for deleted backup copies, the backup copies must be in one of the following states:
 - **Deleted**—The copy is deleted successfully from protection storage, and, if applicable, the agent catalog is deleted successfully from the agent host.
 - **Deleting**—Copy deletion is in progress.
 - **Deletion Failed**—Copy deletion from protection storage is unsuccessful.
 - **Deletion Failed (Agent Catalog)**—The copy is deleted successfully from protection storage, but is not deleted from the agent host.

 **NOTE:** This state is not applicable to virtual machine and Kubernetes backup copies.

 **NOTE:** You cannot export data for backup copies that are in an **Available** state.
5. Click **Download**.
If applicable, the navigation window appears for you to select the location to save the .csv file.
6. Save the .csv file in the desired location and click **Save**.


Remove Exchange, File System, Kubernetes, Block Volume, and SQL backup copies from the PowerProtect Data Manager database

This option enables you to delete the backup copy records from the PowerProtect Data Manager database, but keep the backup copies in protection storage.

About this task

For backup copies that could not be deleted from protection storage, you can remove the backup copies from the PowerProtect Data Manager database. Removing the backup copies from PowerProtect Data Manager does not delete the copies in protection storage.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
2. From the **Assets** window, select the tab for the asset type for which you want to delete copies. If a policy has been assigned, the table lists the assets that have been discovered, along with the associated protection policy.
3. Select a protected asset from the table, and then click **View Copies**. The **Copy Locations** pane identifies where the backups are stored.
4. In the left pane, click  to the right of the icon for the asset. The table in the right pane lists the backup copies.
5. Select one or more backup copies with the **Deletion Failed** status from the table, and then click **Remove from PowerProtect**.
The system displays a warning to confirm that you want to delete the selected backup copies.

8. Click **OK**.

An informational dialog box opens to confirm that the copies are being deleted. To monitor the progress of the operation, click **Go to Jobs**. To view the list of backup copies and their status, click **OK**.

7. Verify that the copies are deleted from the PowerProtect Data Manager database. If the deletion is successful, the deleted copies no longer appear in the table. The backup copies remain in protection storage.

Removing expired backup copies

PowerProtect Data Manager deletes the backup copies of an asset automatically when the retention period of the copy expires. Information about specifying retention periods for a protection policy objective is provided within the user guide for each asset type.

In order for an expired copy to be deleted, the asset must be managed by PowerProtect Data Manager and in one of the following states:

- **Exclusion** – The asset is assigned to an exclusion protection policy.
- **Disabled** – The asset is assigned to a disabled protection policy.
- **Protected** – The asset is assigned to an enabled protection policy.
- **Previously Protected** – The asset has been unassigned from a protection policy and has not yet been reassigned to another policy or assigned to an Exclusion policy.

For an asset assigned to either an exclusion or disabled protection policy, PowerProtect Data Manager deletes the expired backup copies for the asset when the following settings are set to **true**:

- `expiredCopyDeletionEnabledForAssetInExclusionPolicy`
- `expiredCopyDeletionEnabledForAssetInDisabledPolicy`

The expired copy deletion settings for exclusion and disabled protection policies are set to **true** by default. If either setting is set to **false**, PowerProtect Data Manager skips deletion of the expired backup copies. The PowerProtect Data Manager Public REST API documentation provides more information.

Expired copy cleanup occurs at 00:00 AM UTC each day. If a copy deletion fails, a warning alert appears in the audit log under **Alerts > System**.

You can monitor the progress of the expired copy removal job from the **Jobs** window.

Removing assets from PowerProtect Data Manager

PowerProtect Data Manager automatically removes assets if certain conditions are met. However, some assets can be manually removed.

Assets are automatically removed if the following conditions are met:

- The status of the asset is **Deleted**.
- The asset has no backup copies.
- The asset has existed for longer than the value of the asset TTL setting. This is 0 minutes by default, but it can be changed with the REST API. For more information, see PowerProtect Data Manager Public REST API documentation.

NOTE: This value has changed from earlier versions of PowerProtect Data Manager.

The manual removal of assets allows for the following increased control over the process:

- The asset can be removed on demand.
- The status of the asset can be **Not Detected**.
- All protection copies of the asset, including replicated and cloud tiered copies, can be manually removed, followed by the manual removal of the asset.
- All protection copies of the asset can be automatically removed, if this option is selected during manual asset removal from PowerProtect Data Manager.

Remove assets and associated protection copies

In the PowerProtect Data Manager UI, you can manually remove some assets ahead of their scheduled removal, or remove assets that have not been automatically removed.

Prerequisites

- The asset has a status of **Deleted** or **Not Detected**.
- The asset has no protection copies. If copies still exist in the storage system for the asset, you can delete these copies before following the steps in this procedure or select an option to automatically delete the copies when the asset is removed. For information on deleting backup copies, see [Delete backup copies](#).

Steps

1. Select **Infrastructure > Assets**.
2. Select the tab that corresponds to the type of assets that you want to remove. For example, for vCenter virtual machine assets, click **Virtual Machine**.

Assets that are associated with protection copies of this type are listed. By default, only assets with **Available** or **Not Detected** status display. You can also search for assets by name.

3. Select one or more assets from the list, and then click **More Actions > Remove Asset**. The **Remove Assets** dialog displays.

4. Select from one of the following options:

NOTE: All of these options might not display for the selected assets. The available options depend on the protection copy status of the selected assets.

- **Remove assets and associated protection copies**—removes these assets from PowerProtect Data Manager, and automatically removes any protection copies for these assets from storage.
 - **Only remove assets with no associated protection copies**—these assets will not be deleted if PowerProtect Data Manager detects that protection copies for these assets still exist in the storage system.
 - **Mark "Not Detected" assets as "Deleted" but keep associated protection copies**—mark assets with **Not Detected** status as **Deleted** in the PowerProtect Data Manager UI, but retain protection copies for these assets in the storage system. You can view assets marked as **Deleted** from the **Infrastructure > Assets** pane.
5. Click **OK** to confirm the asset removal.

Protecting client assets after a client hostname change

If the hostname of a client is changed, its assets are no longer protected without further action.

When changing a client hostname, you must delete its existing lockbox files and generate new ones. For more information, see the documentation of the relevant application agent.

ifGroup configuration and PowerProtect Data Manager policies

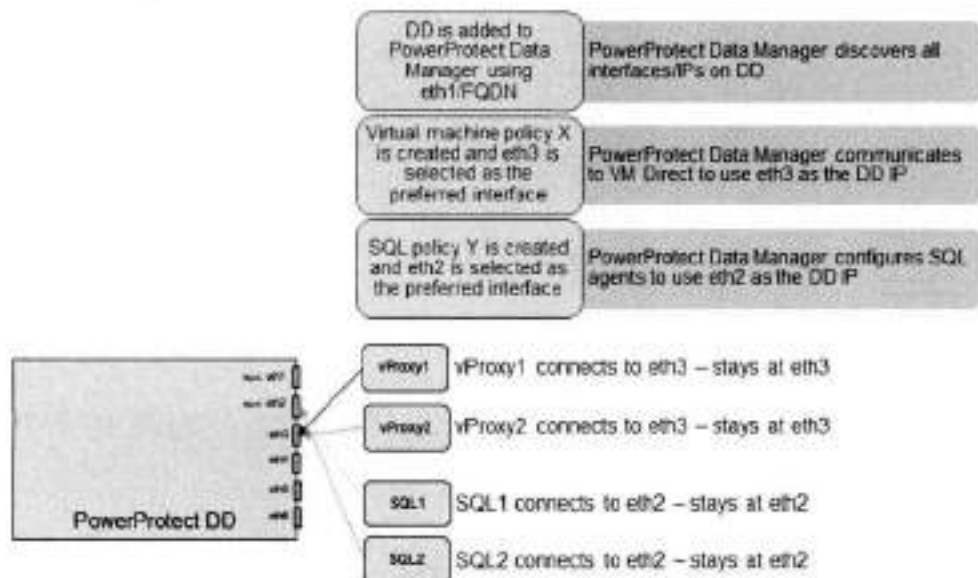
If an `ifGroup` is configured on the DD, the IP address selected in the PowerProtect Data Manager protection policy is only used for the initial connection, and redirection (for example, for load balancing) occurs according to the `ifGroup` setting on the DD. LACP and other failover options on the DD work independently from what is selected in the PowerProtect Data Manager policy.

The following examples and diagrams demonstrate common scenarios in PowerProtect Data Manager when an `ifGroup` is configured on the DD.

PowerProtect Data Manager policy with no `ifGroup`

DD configuration:

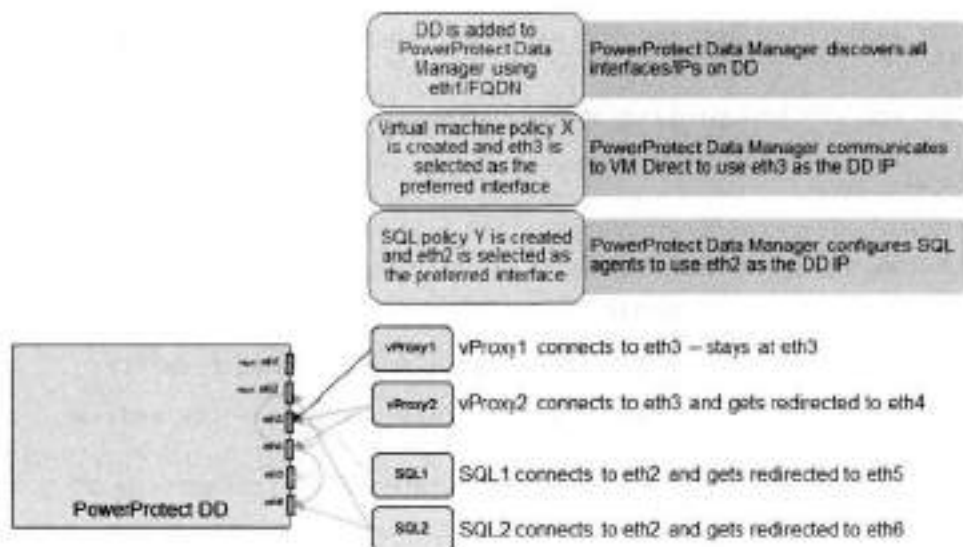
- eth1/eth2 1G
- eth3/eth4/eth5/eth6 10G
- No ifGroup



PowerProtect Data Manager policy with one ifGroup

DD configuration:

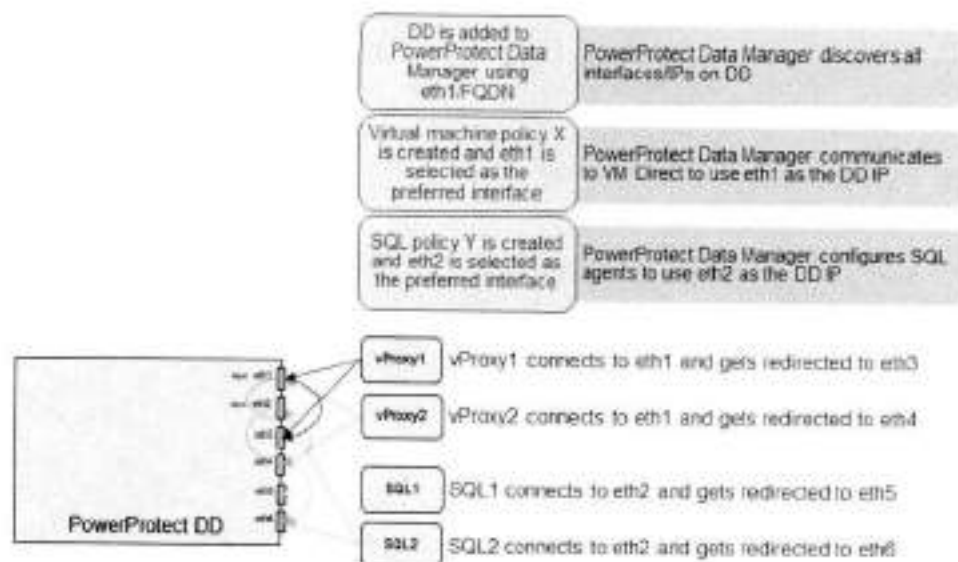
- eth1/eth2 1G
- eth3/eth4/eth5/eth6 10G
- ifGroup * eth3/eth4/eth5/eth6



PowerProtect Data Manager policy with multiple ifGroups

DD configuration:

- eth1/eth2 1G
- eth3/eth4/eth5/eth6 10G
- ifGroup VLAN-VM eth3/eth4
- ifGroup VLAN-SQL eth5/eth6



Troubleshooting failed replication jobs

The following section provides troubleshooting information for when a replication job fails.

Replication to a DD system fails with an authentication error

A replication job might fail with the following error:

```
The backup copies cannot be replicated because the username and password for the source storage system are not valid or cannot be detected.
```

This failure might happen intermittently, while most backup and replication jobs complete without failure and DD systems are successfully discovered.

To resolve this issue, perform the following steps:

1. Collect a DD support bundle and search `/ddvar/log/messages` for `Failed password` or `Invalid user`.

NOTE: For instructions on collecting the support bundle, see the *DDOS Administration Guide*.

2. If the search text is found, you see entries similar to the following:

```
Oct 15 16:36:26 <DD hostname> sshd[25116]: Failed password for sysadmin from <IP address> port 55351
Oct 11 11:18:00 <DD hostname> sshd[31750]: Invalid user <username> from <IP address> port 64425
```

3. Locate the asset-source host using `<IP address>` and correct the credentials that it is using to connect to the DD system.

Replication to a DD system fails with a certificate error

A replication job might fail with an error similar to the following:

```
error = I/O error on POST request for "https://<DD-System>:3009/rest/v1.0/auth":
PKIX path building failed: sun.security.provider.certpath.SunCertPathBuilderException:
unable to find valid certification path to requested target; nested exception is
javax.net.ssl.SSLHandshakeException: PKIX path
```

NOTE: In this example, `<DD_System>` is replaced by the hostname of the replication DD system.

This error indicates that the certificate stored on PowerProtect Data Manager for the DD system has expired, been corrupted, or is missing.

To regenerate the certificate, perform the following steps:

1. If necessary, contact Customer Support to obtain the `ppcp` tool.

NOTE: The `ppcp` tool is not deployed with PowerProtect Data Manager. It must be obtained separately.

2. Log in to the PowerProtect Data Manager server console using administrator credentials and change to the root user.
3. Change directories to the location of the `ppcp` tool.
4. Run the command `./ppcp rest --uri certificates --params "host=<DD_System>&port=3009&type=HOST"`, replacing `<DD_System>` with the hostname of the DD system.
5. In the command output similar to the following, record the value of `id` inside quotation marks immediately after `"id"`:

```
"fingerprint":*473303EFF3EFE6D6AAC2D76F1FB94561B12A321F", "host": "<DD_System>", "id": "2G1vYmstcGRiMDAwMS1wLmQuaW50LmluZnJhLjdhbmRpLmNvLmpwOjMwMDk6aG9zdA==", "issuerName": "CN=<DD_System>, OU=Root CA, O=Valued Datadomain Customer, L=Santa Clara, ST=CA, C=US", "notValidAfter": "Sun Dec 07 15:28:39 JST 2025", "notValidBefore": "Mon Dec 09 00:28:39 JST 2019", "port": "3009", "state": "UNKNOWN", "subjectName": "CN=<DD_System>, O=Valued DataDomain customer, OU=Root Certificate, ST=CA, C=US", "type": "HOST"
```

NOTE: In this example, the value of `id` is

```
2G1vYmstcGRiMDAwMS1wLmQuaW50LmluZnJhLjdhbmRpLmNvLmpwOjMwMDk6aG9zdA==
```

6. Run the following command to change the state of the certificate from `UNKNOWN` to `ACCEPTED`:
 - Use the value of `id` recorded in step 5. The value of `id` is entered immediately after `certificates/`.
 - With the exception of changing `UNKNOWN` to `ACCEPTED`, copy the output of the command from step 4 between `{` and `}`.
 - Replace `<DD_System>` with the hostname of the DD system.
 - This command is entered on a single line of text.

```
./ppcp rest --method PUT --uri certificates/  
2G1vYmstcGRiMDAwMS1wLmQuaW50LmluZnJhLjdhbmRpLmNvLmpwOjMwMDk6aG9zdA== --data  
{ "fingerprint":*473303EFF3EFE6D6AAC2D76F1FB94561B12A321F", "host": "dmobk-pdb0001-  
p.d.int.infra.7andi.co.jp", "id": "2G1vYmstcGRiMDAwMS1wLmQuaW50LmluZnJhLjdhbmRpLmNvLmpwOjMwMDk6aG9zdA==", "issuerName": "CN=dmobk-pdb0001-p.d.int.infra.7andi.co.jp, OU=Root  
CA, O=Valued Datadomain Customer, L=Santa Clara, ST=CA, C=US", "notValidAfter": "Sun  
Dec 07 15:28:39 JST 2025", "notValidBefore": "Mon Dec 09 00:28:39 JST  
2019", "port": "3009", "state": "ACCEPTED", "subjectName": "CN=dmobk-pdb0001-  
p.d.int.infra.7andi.co.jp, O=Valued DataDomain customer, OU=Host Certificate, ST=CA,  
C=US", "type": "HOST" }
```

Restoring Data and Assets

Topics:

- View backup copies available for restore
- Restoring a protection policy
- Restore the PowerProtect Data Manager server
- Restore Cloud Tier backups to protection storage

View backup copies available for restore

When a protection policy is successfully backed up, PowerProtect Data Manager displays details such as the name of the storage system containing the asset backup, location, the creation and expiry date, and the size. To view a backup summary:

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Assets** or **Restore > Assets**.
2. Select the tab that corresponds to the type of assets that you want to view. For example, for vCenter virtual machine assets, click **Virtual Machine**.

Assets that are associated with protection copies of this type are listed. By default, only assets with **Available** or **Not Detected** status display. You can also search for assets by name.

For virtual machines, you can also click the **File Search** button to search on specific criteria.

NOTE: In the **Restore > Assets** window, only tabs for asset types supported for restore within PowerProtect Data Manager display. Supported asset types include the following:

- **Virtual Machine**
- **File System**
- **Storage Group**
- **Kubernetes**
- **Network Attached Storage (NAS)**
- **Oracle**
- **SQL**
- **Block Volumes**

3. To view more details, select an asset and click **View copies**.

The copy map consists of the root node and its child nodes. The root node in the left pane represents an asset, and information about copy locations appears in the right pane. The child nodes represent storage systems.

When you click a child node, the right pane displays the following information:

- Storage system where the copy is stored.
- The number of copies
- Details of each copy, including the time that each copy was created, the consistency level, the size of the copy, the backup type, the copy status, and the retention time.
- The indexing status of each copy at the time of copy creation:
 - **Success** indicates that all files or disks are successfully indexed.
 - **Partial Success** indicates that only some disks or files are indexed and might return partial results on file search.
 - **Failed** indicates that all files or disks are not indexed.
 - **In Progress** indicates that the indexing job is in progress.

If indexing has not been configured for a backup copy, or if global expiration has been configured and indexed disks or files have been deleted before the backup copy expiration date, the **File Indexing** column displays **N/A**.

The indexing status updates periodically which enables you to view the latest status.

- For virtual machine backups, a **Disk Excluded** column enables you to view any virtual disks (VMDKs) that were excluded from the backup.

Restoring a protection policy

You can use the PowerProtect Data Manager user interface to perform centralized and self-service restores of protection policy backups for any of the following asset types. For more information, see the appropriate publication.

Table 27. Protection-policy asset types

Asset type	Publication
File system data	<i>PowerProtect Data Manager File System User Guide</i>
Kubernetes cluster namespaces and PVCs	<i>PowerProtect Data Manager Kubernetes User Guide</i>
Microsoft Exchange Server databases	<i>PowerProtect Data Manager Microsoft Exchange Server User Guide</i>
Microsoft SQL Server databases	<i>PowerProtect Data Manager Microsoft SQL Server User Guide</i>
Network Attached Storage (NAS) share and appliance data	<i>PowerProtect Data Manager Network-Attached Storage User Guide</i>
Oracle RMAN databases	<i>PowerProtect Data Manager Oracle RMAN User Guide</i>
SAP HANA databases	<i>PowerProtect Data Manager SAP HANA User Guide</i>
Storage Direct VMAX storage groups	<i>PowerProtect Data Manager Storage Direct User Guide</i>
Virtual machines	<i>PowerProtect Data Manager Virtual Machine User Guide</i>
Block volumes	<i>PowerProtect Data Manager Storage Array User Guide</i>

Restore the PowerProtect Data Manager server

You can restore PowerProtect Data Manager server persisted data as a new instance using any of the backups. Only the Administrator role can carry out the restore.

Prerequisites

Ensure that:

- The PowerProtect Data Manager version that is deployed on your system and the backups you are using for the restore match.
- The network configuration is the same on the newly deployed PowerProtect Data Manager system as on the failed instance that you are restoring.

Steps

1. Deploy the PowerProtect Data Manager OVA and power it on.
2. Select **Restore Backup**.

To delay jobs defined by your protection policies until otherwise specified, select **After restore, keep the product in recovery mode so that scheduled workflows are not triggered**. When selected, after restore the system enters recovery maintenance mode. During recovery maintenance mode:

- All jobs defined by your protection policies that modify the backup storage (for example, backup creation, backup deletion, and PowerProtect Data Manager Server DR jobs) are not triggered.
- All operations that write to the backup storage are disabled.
- A system alert is displayed in PowerProtect Data Manager.



To enable automatically scheduled operations and user operations that write to the backup storage, click **Return to full Operational mode** in the alert.

3. Specify the following storage information:
 - a. DD system IP where the recovery backups are stored.
 - b. DD NSF Export Path where the recovery backups are stored.
 - c. Click **Connect**.
4. Select the PowerProtect Data Manager instance that you would like to restore, and then click **OK**.
5. Select the backup file that you would like to use for recovery, and then click **Recover**.
6. Specify the lockbox passphrase associated with the backup, and start the recovery.
This step initiates the recovery and display the progress status. The recovery process can take approximately eight minutes before the URI is redirected to the PowerProtect Data Manager login.

Results

The PowerProtect Data Manager server is recovered.

Next steps

After a successful recovery:

- The time zone of the PowerProtect Data Manager instance is set to the same as that of the backup.
- All preloaded accounts are reset to default passwords, as described in the *PowerProtect Data Manager Security Configuration Guide*. The preloaded UI administrator account is an exception and retains its password. Change the passwords for all preloaded accounts as soon as possible.

Restore Cloud Tier backups to protection storage

Once a Cloud Tier backup is recalled, restore operations of these backups are identical to normal restore operations.

The PowerProtect Data Manager software recalls a copy of the backup from the Cloud Unit to the local (active) tier of protection storage, which then allows you to perform a restore of the backup from the active tier to the client. The status appears as **Cloud**, and changes to **Local Recalled** after cloud recall completes. After the restore, the backup copy is removed from Cloud Tier, and is stored on the active tier of protection storage for a minimum of 14 days, after which the backup may be returned to the cloud depending on your protection policy.

Recall and restore from Cloud Tier


Perform the following steps to recall a backup on Cloud Tier to the active tier on protection storage and restore this backup.

Prerequisites

NOTE: When a backup is recalled from Cloud Tier to the active tier, the copy is removed from Cloud Tier.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
2. On the **Assets** window, select the tab that contains the asset you want to recall from Cloud Tier, and then click **View Copies**.
3. Click **DD**, and then select from one of the available copies that appear in the table.

4. Click **Recall**.
The **Recall from Cloud** dialog box appears.
5. In the **Retain until** box, specify how long you want to keep the copy on the active tier, and then click **OK**.
6. Go to the **Jobs** window to monitor the recall operation.
When the copy has been moved successfully, the **Location** changes from Cloud to Local.
7. Select **Restore > Assets**, and then select the tab that contains the recalled asset.
8. Select the recalled asset, and then click **Restore**.
 **NOTE:** If you are unsure whether the asset has been recalled, click **View Copies** and select **DD** to view the available backup copies. If the asset backup is a recalled copy, the Status column indicates **Local Recalled**.
9. Select the recalled copy to re-tier the copy to the active tier.

Preparing for and Recovering From a Disaster

Topics:

- About server disaster recovery
- System recovery for server DR
- Quick recovery for server DR
- Overview of PowerProtect Data Manager Cloud Disaster Recovery

About server disaster recovery

The PowerProtect Data Manager system protection service enables you to protect the persistent data of a PowerProtect Data Manager system from catastrophic loss by creating a series of server disaster recovery (DR) backups.

Preparing for server DR requires the consideration of two scenarios: loss of the PowerProtect Data Manager server and loss of the entire site. Some of the information that you record during the server DR configuration process may only apply to one scenario or the other. As a best practice, you should gather and record all applicable information for both scenarios.

PowerProtect Data Manager supports three methods of server DR:

System recovery

System recovery creates point-in-time snapshots of the PowerProtect Data Manager server in protection storage. During a DR activity, recover the server from protection storage and then restore protected assets.

System recovery for server DR provides more information.

Quick recovery

Quick recovery makes a remote PowerProtect Data Manager replication destination aware of replicated backups and enables the recovery view. During a DR activity, you can restore assets from these replicated backups at the destination without first restoring the source server.

Quick recovery for server DR provides more information.

Cloud Disaster Recovery

Cloud DR enables you to restore to a DR site in a supported public cloud environment. During a DR activity, restore virtual machines to a Cloud DR server and recover those workloads in the cloud.

Overview of PowerProtect Data Manager Cloud Disaster Recovery provides more information.

Differences between server DR methods

The following table highlights the differences between the three server DR methods.

Table 28. Server DR comparison

Criteria	System recovery	Quick recovery	Cloud DR
Requires another running PowerProtect Data Manager server	No ^a	Yes	No
Requires additional configuration after setup	Optional ^b	No	No
Requires configuration outside of the PowerProtect Data Manager UI during recovery	Yes	No	No
Preserves backup workflows	Yes	No	No
Supports server DR replication	Yes	Automatic	Automatic
Recovery time objective (RTO) for backup infrastructure	>1 hour ^a	N/A	N/A

- a. Optionally, you can configure a second server and leave this server unconfigured to decrease the RTO for system recovery. However, the RTO for system recovery cannot match the RTO for quick recovery or Cloud DR.
- b. Configuration of server DR replication.

System recovery for server DR

The system recovery process creates periodic backups of a PowerProtect Data Manager server, from which you can restore the server after a disaster. Each backup is considered a full backup although it is created in an incremental manner.

System recovery backups include persistent data such as the lockbox and the PowerProtect Data Manager databases. The backup operation quiesces the server and creates a point-in-time snapshot of the databases. This quiescent state limits user functionality. After the snapshot completes and while PowerProtect Data Manager copies the snapshot to protection storage, the server restores full user functionality. System recovery backups also include File Search indexes and other component DR backups.

The system protection service enables you to manage the frequency and retention of an automated server DR backup. You can also perform manual backups. However, the system protection service does not manage the retention of manual backups and you must delete any outdated manual backups yourself. Manage PowerProtect Data Manager server DR backups provides instructions.

You can select one protection storage system as a server DR backup target and one protection storage system as a replication target. Replication provides an extra layer of protection for server DR backups. Manually configure server DR backups provides instructions for configuring server DR replication, while Recover PowerProtect Data Manager from server DR backups contains instructions for restoring from a replica.

Since only one backup target and one replication target are supported at a time, when you specify a new protection storage system, you overwrite the existing selection. If you have more protection storage systems, you can change which protection storage system holds the server DR backup or receives the replica.

PowerProtect Data Manager server DR replication is independent of any legacy methods, such as MTree replication on an individual DD system. Backups and configuration from legacy methods are not detected or migrated.

Server DR protection storage types

PowerProtect Data Manager supports two types of protection storage for server DR: NFS and DD Boost.

DD Boost is the recommended storage type for PowerProtect Data Manager server DR. NFS is the legacy storage type for PowerProtect Data Manager server DR.

Updating the PowerProtect Data Manager server does not automatically change the storage type. Instead, select the appropriate storage type and manually configure server DR backups. Do not alternate between storage types.

Switching from NFS to DD Boost creates new server DR backups, rather than migrating existing backups. The previous NFS backups are no longer visible in the list of DR backups. However, you can still recover from older NFS server DR backups even after switching to DD Boost, should you experience a disaster before the initial DD Boost system backup completes.

DD Boost

DD Boost provides security and efficiency advantages over NFS, including password-protected authentication. When you use DD Boost, PowerProtect Data Manager creates and manages a storage unit on the DD system and a corresponding user account.

- The storage unit and user account name are based on the PowerProtect Data Manager hostname. For example, `sysDR_<hostname>`.
- The DD Boost user password is based on the PowerProtect Data Manager predefined administrator account (`admin`) password.

NOTE: The password is based on the `admin` account even if you use other accounts with the Administrator role, such as external identity provider users, to administer PowerProtect Data Manager.

Changes to the PowerProtect Data Manager predefined administrator password prompt corresponding updates to the DD Boost user password. If you configured server DR replication, password changes also prompt corresponding updates to the credentials on the replication target. Recovery from server DR backups requires the PowerProtect Data Manager predefined administrator password. If you do not know this password, contact Customer Support.

If you plan to use DD Boost, add the DD system as protection storage before you configure server DR. Protection storage provides instructions.

The DD Boost storage type allows for automatic server DR configuration. Automatic server DR provides more information. Only the DD Boost storage type supports server DR replication.

NFS

To store backups over NFS, you must configure and assign a private storage unit for the PowerProtect Data Manager system. Then, prepare the DD recovery target by creating an NFS export. With the DD system address and the NFS export path, you can configure PowerProtect Data Manager to perform server DR backups.

NFS storage is deprecated in favor of DD Boost.

Automatic server DR

New deployments of PowerProtect Data Manager automatically configure and enable server DR with minimal input. This process ensures that the server is protected as soon as you add protection storage.

Automatic server DR detects when you first add a protection storage system. The automatic configuration mechanism uses the recommended DD Boost storage type and default settings to create a managed storage unit for server DR. This process generates server DR jobs that you can track through the **Jobs** page.

Automatic configuration selects the first protection storage system that you add to PowerProtect Data Manager. However, you can configure server DR to change the target to another protection storage system or enable replication. *Manually configure server DR backups* provides instructions. Manual configuration of the backup target is not recommended unless you must target a different protection storage system.

If automatic server DR fails, *Manually configure server DR backups* provides an alternate method to configure server DR. The job details provide information that you can use to troubleshoot the configuration process.

Prepare the DD system recovery target (NFS)

If you plan to use NFS for system backup storage, configure the NFS export on the DD target system and select the required permissions. Configuring PowerProtect Data Manager for backup and recovery requires this NFS export path.


About this task

NOTE: NFS is the legacy storage type for PowerProtect Data Manager server DR.

Steps

1. Use a web browser to log in to DD System Manager as the system administrator user.
2. On the **Summary** tab in the **Protocols** pane, select **NFS Exports > Create Export**.

- In the **Create NFS Export** window, provide the following information, and then click **OK**.
 - Export Name**—the name of the DD MTree.
 - Directory Path**—the full directory path for DD MTree that you created. Ensure that you use the same name for the directory.

 **NOTE:** For an external DD system, specify a path similar to the following, `/data/coll/<path>`, where `<path>` is the MTree that stores the system backups.
- Add PowerProtect Data Manager by hostname or IP address to the NFS client list.
To configure DR protection for an existing Search cluster, add the IP address or hostname of the Search cluster to the NFS client list.
- Ensure that the **Current Selection** list includes `no_root_squash`, which is required for permission for PowerProtect Data Manager to change the directory structure on the NFS share.
- When the progress message indicates that the save operation is complete, click **Close**.

Manually configure server DR backups


For new deployments, PowerProtect Data Manager automatically configures and enables server DR. However, you can manually configure DR protection for the PowerProtect Data Manager system and the system metadata.

Prerequisites

If you plan to use NFS for protection storage, prepare the target DD system as described in Prepare the DD system recovery target (NFS).

If you plan to use DD Boost for protection storage, add the DD system as protection storage. Protection storage provides instructions. If you plan to replicate server DR backups, the replication target must be a different protection storage system.

Steps

- Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.
- Click  select **Disaster Recovery**, and then click **Configuration**.
- Select **Enable backup**.
- For DD Boost, configure the backup with the following attributes:
 - For **Protocol**, select **DDBoost**.
 - From the **PowerProtect DD System** drop-down list, select a backup destination from the list of existing protection storage systems, or select **Add** to add a system and complete the details in the **Add Storage** window.
For initial DR configuration, the **Storage Unit** field is empty. If DR was already configured, the **Storage Unit** field displays the name of the storage unit that holds server DR backups.
- For NFS, configure the backup with the following attributes:
 - For **Protocol**, select **NFS**.
 - In the **PowerProtect DD System** field, type the IP address or hostname of the DD system for the backup.
 - In the **NFS Export Path** field, type the NFS path where server DR backups are stored on the target DD system.
- Configure the backup frequency and duration:
 - Type an interval between server DR backups, in hours.
This setting controls backup frequency, and the allowed values are 1 to 24 hours.
 - Type the number of days for which PowerProtect Data Manager should retain server DR backups.
The allowed values are 2 to 30 days.
- To enable server DR replication:
 - Check **Enable Replication**.
 - From the **Replicate Backup To** drop-down list, select a target from the list of existing protection storage systems, or select **Add** to add a system and complete the details in the **Add Storage** window.
The replication target cannot be the backup destination.
The replication frequency and retention time are the same as for the backup.
- Click **Save**.

Results

For DD Boost, PowerProtect Data Manager creates system jobs to prepare the new storage unit and to configure the server DR protection policy.

For both storage types, PowerProtect Data Manager creates a system job for the first server DR backup.

If you configured replication, PowerProtect Data Manager creates a DD Boost user and storage unit on the destination. Server DR backups begin replicating according to the protection schedule.

Next steps

Verify that the system jobs succeed.

Record settings for server DR

Plan for DR by recording vital information. In the event of a major outage, you will need this information to recover your systems. Some items are only required for particular DR scenarios. Record the following information on a local drive outside PowerProtect Data Manager:

Steps

1. If PowerProtect Data Manager is deployed to vSphere, record the port groups:
 - a. Log in to the vSphere client.
 - b. Right-click the appliance name and select **Edit Settings**.
 - c. Record the port group settings that are assigned to PowerProtect Data Manager.
This information is useful when restoring to the same VMware environment.
2. Record the PowerProtect Data Manager FQDN.
3. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.
4. Record the PowerProtect Data Manager version and build numbers.
Customer Support can provide this information, which is not mandatory.
5. Click  select **Disaster Recovery**, and then click **Configuration**.
6. Record whether server DR storage uses NFS or DD Boost.
7. Record the protection storage system IP address or FQDN.
8. If you configured server DR replication, record the FQDN for the replication target.
9. If you use NFS for server DR storage, record the NFS export path.
10. If you use DD Boost for server DR storage, perform the following substeps:
 - a. Connect to the PowerProtect Data Manager console and change to the root user.
 - b. Change directory:
`cd /usr/local/bras/puppet/scripts`
 - c. Obtain and record the server DR DD Boost credentials:
`./get_sdr_config_credential.py SysDR_$(hostname -s)`
 - d. Connect to the protection storage system console.
 - e. Obtain and record the user ID (UID) for the server DR DD Boost user:
`user show list`

Results

Table 29. Recorded DR settings

System	Setting or Property	Example	Recorded Value
PowerProtect Data Manager	Version and build	19.14	
	FQDN	server1.example.com	
	Backup protocol	NFS or DD Boost	
Server DR replica	FQDN	dd-replica.example.com	

Table 29. Recorded DR settings (continued)

System	Setting or Property	Example	Recorded Value
Protection storage system	FQDN	dd.example.com	
	NFS export path	N/A	
	DD Boost username	SysDR_server1	
	DD Boost password	zD0_56c-b4e-ad4-dbb-	
	DD Boost UID	501	

Manage PowerProtect Data Manager server DR backups

View PowerProtect Data Manager server DR backups and perform manual backups.

About this task


For DR backups, PowerProtect Data Manager supports a default retention period of 7 days plus the last 3 hourly backup copies for the current day. You can change the frequency and retention of DR backups from the **Disaster Recovery > Configuration** tab.

The system protection service automatically deletes scheduled backups according to the configured retention policy.

You can manually delete all backups except for the most recent backup marked as **FULL** and the most recent backup marked as **PARTIAL**.

Steps

1. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.

2. Click  select **Disaster Recovery**, and then click **Manage Backups**.

3. To perform a manual backup:

a. Click **Backup Now**.

The **Enter a name for your backup** dialog appears.

b. [Optional] Type a name for your backup.

You can leave the backup name blank, and PowerProtect Data Manager provides a name for the backup by using the naming convention `DiserDR-`. If you provide a name with the convention that PowerProtect Data Manager uses for scheduled backups, which is `SystemDR`, PowerProtect Data Manager displays an error.

c. Click **Start Backup**.

The backup appears as an entry in the table. To view details for the backup, click **>**.

If the Search Engine is deployed, PowerProtect Data Manager also backs up the Search Engine. The backup details provide the status of the Search Engine backup.

To monitor the status of the backup, select **Jobs > Protection** and look for a job with the name **Protect the server datastore**.

4. To delete a backup:

a. Select a backup from the list.

b. Click  for that row.

The system displays a warning to confirm you want to delete the backup. Click **Yes** to proceed.

5. Click **Cancel**.

Recover PowerProtect Data Manager from server DR backups

You can recover PowerProtect Data Manager from a server DR backup on a protection storage system.

Prerequisites

- Only the Administrator role can carry out the recovery.

- Ensure that all the information listed in Record settings for server DR is available.
- Ensure that the FQDN of the PowerProtect Data Manager is the same as the host name.
- To restore data from NFS, ensure that you have set up the recovery target system. See Prepare the DD system recovery target (NFS).
- To restore data from DD Boost, ensure that you have the current password for the PowerProtect Data Manager UI predefined administrator account. If you do not know this password, contact Customer Support.
- To restore data from a server DR replica, ensure that you have the IP address or FQDN for the replication target, the PowerProtect Data Manager hostname, and the current password for the PowerProtect Data Manager UI predefined administrator account.
- If the Search Engine or reporting engine nodes from the previous PowerProtect Data Manager deployment are still hosted on the vCenter server, delete the Search Engine and reporting engine nodes from the vCenter server before you recover the PowerProtect Data Manager system. The recovery process redeploys the Search Engine and reporting engine nodes as part of the recovery operation.
- The recovery process does not automatically redeploy protection engines. After recovery, redeploy the protection engines.

About this task

When a primary PowerProtect Data Manager system fails because of a major event, deploy a new PowerProtect Data Manager system and recover the backup from the external DD system.

NOTE: If the recovery system has a different FQDN, see Troubleshoot recovery of PowerProtect Data Manager.

If a Search Engine is present in the recovery backup when you recover the PowerProtect Data Manager system, the Search Engine is automatically recovered.

Steps

1. Deploy a new PowerProtect Data Manager virtual appliance.
The *PowerProtect Data Manager Deployment Guide* for the appropriate platform provides instructions.
2. From a host that has network access to the virtual appliance, use the latest version of Google Chrome to connect to the appliance:
`https://<appliance_hostname>`
NOTE: You can specify the hostname or the IP address of the appliance.
3. On the **Install** window under **Welcome**, select **Restore Backup**.
4. Select **After restore, keep the product in recovery mode so that scheduled workflows are not triggered**.
Recovery mode provides more information.
5. To restore data from NFS:
 - a. For **Protocol**, select **NFS**.
 - b. Under **Select File**, enter the DD System and NFS Export Path where the backup is located, and then click **Connect**.
A list of the available recovery backups appears.
6. To restore data from DD Boost:
 - a. For **Protocol**, select **DDBoost**.
 - b. Type the hostname or IP address for the protection storage system that stores server DR backups.
 - c. If the hostname is not already populated, type the hostname for the original PowerProtect Data Manager system.
 - d. To restore from a server DR replica, append **/R** to the hostname.
For example, `system1.example.com/R`.
 - e. Type the password for the predefined administrator account (`admin`) of the original PowerProtect Data Manager.
 - f. Click **Connect**.
A list of the available recovery backups appears. If restoring data from a replica, the list of backups includes those on the replica.
7. Select the backup from which to recover the system, and then click **Recover**.
The recovery starts. Recovery can take a few minutes.
NOTE: There is no busy indicator during the recovery process. The current recovery state can be monitored from the text displayed in the recovery window.

Results

When recovery is complete, the PowerProtect Data Manager login page appears.

The time zone of the PowerProtect Data Manager instance is set to that of the backup.

If restoring from a replica, the replication target protection storage system is configured as the new server DR backup target.

All preloaded accounts are reset to default passwords, as described in the *PowerProtect Data Manager Security Configuration Guide*. The preloaded UI administrator account is an exception and retains its password. Change the passwords for all preloaded accounts as soon as possible.

NOTE: Backup copies that were created after the Server DR recovery backup used in the recovery process are discovered after the server is recovered. However, any backup copies that had replication or cloud tier copies before the recovery operation are replicated or cloud tiered during the next manual or scheduled job.

Recovery mode

If you select **After restore, keep the product in recovery mode so that scheduled workflows are not triggered** during deployment, PowerProtect Data Manager enables recovery mode.

With recovery mode active, when you log in to PowerProtect Data Manager:

- A red banner appears at the top of the PowerProtect Data Manager UI. The banner indicates that the PowerProtect Data Manager system is operational but scheduled workflows are disabled.
- All jobs defined by your protection policies that modify the backup storage (for example, backup creation, backup deletion, and PowerProtect Data Manager Server DR jobs) are not triggered.
- All operations that write to the backup storage are disabled.

To return PowerProtect Data Manager to full operational mode and enable scheduled workflows, click **Return to full operational mode**.

Recover the Search Engine from a DR backup

PowerProtect Data Manager automatically restores the Search cluster after disaster recovery of the PowerProtect Data Manager system is complete. If the PowerProtect Data Manager system could not restore the Search cluster automatically, use the steps in this procedure to restore only the Search cluster through the REST API. Recovery of a Search cluster must be performed on an operational PowerProtect Data Manager system. Only the Administrator role can restore the Search cluster.

Prerequisites

Obtain the name of the Search cluster backup from **System Settings > Disaster Recovery > Manage Backups**.

About this task

Use the backup manifest file to create a new text document that will be used to issue a POST command with the REST API:

CAUTION: Do not edit the manifest file itself.

Steps

1. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.
Use the same credentials that you used before PowerProtect Data Manager was restored.
2. Connect to the PowerProtect Data Manager console as an admin user.
3. Change directories to `/data01/server_backups/<PowerProtect Data Manager Hostname>_<NodeID>` to locate the backup manifest file.

Normally, there is only a single subdirectory in `/data01/server_backups`, so change to that subdirectory. However, if there is more than one subdirectory and you don't know which `<NodeID>` is the correct one, perform the following substeps:

- a. From `/data01/server_backups`, run the following commands, changing the username and password as required:

```
TOKEN=$(curl -X POST https://localhost:8443/api/v2/login -k -d '{ "username":  
"admin","password": "admin_password" }' --header "Content-Type: application/json" |  
python3 -c "import sys, json; print(json.load(sys.stdin)['access_token'])")
```

```
curl -X GET https://localhost:8443/api/v2/nodes -k --header "Content-Type:  
application/json" --header "Authorization:Bearer $TOKEN"
```


- b. Run the command `grep -Rnwa -e '<Name>' --include=*.manifest`.
4. Copy the manifest file to a temporary file.
5. Open the temporary file.
6. Review the following example, and make the changes documented by the `//` comment entries.

NOTE: The `//` comment entries displayed here do not exist in the temporary file itself. These comment entries are displayed here only as a guide.

```
{
  "id": "ca8cbb13-6f3d-4ac5-87e5-de47a634379f",
  "jobId": "990b4ea7-c0e4-4069-8dd5-7d0e084370fc", // DELETE LINE
  "creationTime": "2022-08-25T19:38:54.622275+0000",
  "lastUpdated": "2022-08-25T19:40:18.165497Z", // DELETE LINE
  "elapsedSeconds": 11,
  "sequenceNumber": 2,
  "state": "Successful", // DELETE LINE
  "version": "19.12.0-1-SNAPSHOT", // DELETE LINE
  "hostname": "ldpdb141.hop.lab.emc.com", // DELETE LINE
  "name": "mercijTestDr", // DELETE LINE
  "nodeId": "a8d2df8e-5c3e-4160-87d4-32b9bfe6c283", // DELETE LINE
  "sizeInBytes": 29759075,
  "consistency": "CRASH CONSISTENT", // DELETE LINE
  "checksum": "bbd97a04f296a8ed116e4a9272982d9e9411f3d0cf50dea131d5c2cd4ce224f8", //
DELETE LINE
  "backupConsistencyType": "FULL", // DELETE LINE
  "esSnapshotState": "UNKNOWN", // DELETE LINE
  "backupTriggerSource": "USER", // DELETE LINE
  "configType": "standalone", // DELETE LINE
  "deployedPlatform": "vmware", // DELETE LINE
  "replicationTargets": [], // DELETE LINE
  "repositoryFileSystem": "BOOST_FILE_SYSTEM", // DELETE LINE
  "ddHostname": "ldpdg251.hop.lab.emc.com", // DELETE LINE and add line
  "recover": true,
  "Components": [ // change Components to components with lower case c
    { // DELETE WHOLE PPDM COMPONENT LEAVING ONLY SEARCHCLUSTER
      "name": "PPDM",
      "id": "ca7cbb13-6f3d-4ac5-87e5-de47a634379f",
      "lastActivityId": "2bdbe7a8-7c57-446d-b072-ad8081e2953d",
      "version": "v2",
      "backupPath": "ldpdg251.hop.lab.emc.com:SysDR_ldpdb141/
ldpdb141_a8d2df8e-5c3e-4160-87d4-32b9bfe6c283/PPDM",
      "backupStatus": "SUCCESSFUL",
      "backupsEnabled": true,
      "errorResults": []
    }, // STOP DELETING HERE
    {
      "name": "SearchCluster",
      "id": "ca7cbb13-6f3d-4ac5-87e5-de47a634379f",
      "lastActivityId": "198a93b1-7382-474b-89c8-c7b6b0ab4987",
      "version": "v2",
      "backupPath": "ldpdg251.hop.lab.emc.com:SysDR_ldpdb141/
ldpdb141_a8d2df8e-5c3e-4160-87d4-32b9bfe6c283/SearchCluster",
      "backupStatus": "SUCCESSFUL",
      "backupsEnabled": true, // DELETE TRAILING COMMA
      "errorResults": [] // DELETE LINE
    }
  ]
}
```

In summary:

- remove all lines with the `// DELETE LINE` comment entry displayed here
- add `recover: true`
- change `Components` to `components`
- remove all listed component blocks except for `Search Cluster`
- remove the trailing comma from `"backupsEnabled": true,`

The result of these changes should look similar to the following:

```
{
  "id": "ca8cbb13-6f3d-4ac5-87e5-de47a634379f",
  "creationTime": "2022-08-25T19:38:54.622275+0000",
  "elapsedSeconds": 11,
  "sequenceNumber": 2,
  "sizeInBytes": 29759075,
  "recover": true,
  "components": [
    {
      "name": "SearchCluster",
      "id": "ca7cbb13-6f3d-4ac5-87e5-de47a634379f",
      "lastActivityId": "198a93b1-7382-474b-89c8-c7b6b0ab4987",
      "version": "v2",
      "backupPath": "ldpdg251.hop.lab.emc.com:SysDR_ldpdb141/
ldpdb141_a8d2df8e-5c3e-4160-87d4-32b9bfe6c283/SearchCluster",
      "backupStatus": "SUCCESSFUL",
      "backupsEnabled": true
    }
  ]
}
```

- Copy the value of the text inside the quotation marks that follow "id":.
This value replaces the variable <backupID> used in step 11. In this example, <backupID> is ca8cbb13-6f3d-4ac5-87e5-de47a634379f.
- Remove all carriage returns from the temporary file, so that all the text is on a single line.
- Copy all of the text from the temporary file.
This value replaces the variable <manifestText> used in step 11.
- Run the following command, changing the username and password credentials as required:

NOTE: Even if you ran this command in step 5.a, run it again. The validity of the value of TOKEN is time sensitive.

```
TOKEN=$(curl -X POST https://localhost:8443/api/v2/login -k -d '{"username":
"admin","password": "admin_password"}' --header "Content-Type: application/json" |
python3 -c "import sys, json; print(json.load(sys.stdin)['access_token'])")
```

- Run the following command:

```
curl -X PUT 'https://localhost:8443/api/v2/server-disaster-recovery-backups/
<backupID>' --header "Authorization: Bearer $TOKEN" --header 'Content-Type:
application/json' -k -d '<manifestText>'
```

- Replace <backupID> with the value obtained in step 7.
 - Replace <manifestText> with all of the text obtained in step 9.
- To monitor the status of the restore process in the PowerProtect Data Manager UI, select **Jobs > System Jobs** and look for a job with the description **Restoring backup Search Node**.

Next steps

Delete the temporary file created in step 4.

Change the IP address or hostname of a DD system

You can change the IP address or hostname of a DD system without affecting server DR.

About this task

Before changing the IP address or hostname of a DD system, perform the following steps.

- NOTE:** If you have changed the IP address or hostname of a DD system without following these steps, you can recover DR functionality. For more information, see [Recover from a changed DD system IP address or hostname](#).

Steps

1. Disable Server DR backups.
2. Log in to the PowerProtect Data Manager server by using SSH.
3. Run the following command:

```
sudo umount /data01/server_backups
```
4. For each Search Engine node, perform the following substeps:
 - a. Log in to the Search Engine node by using SSH.
 - b. Run the following command:

```
sudo umount /mnt/PPDM_Snapshots
```
5. Remove the DD system from PowerProtect Data Manager:
 - a. From the PowerProtect Data Manager UI, select **Infrastructure > Storage**.
 - b. Select the DD system to remove.
 - c. Click **Delete**.
6. Change the IP address or hostname of the DD system.
7. Add the DD system back to PowerProtect Data Manager.
8. Enable server DR backups.

Recover from a changed DD system IP address or hostname

If you changed the IP address or hostname of a DD system without following the supported procedure, you can recover your server DR functionality.

About this task

Steps

1. Disable Server DR backups.
2. Log in to the PowerProtect Data Manager server by using SSH.
3. Run the following command:

```
ps aux | grep /data01/server_backups | grep boostfs
```

Make a note of the process ID next to the `boostfs` entry in the command output.
4. Run the following command, replacing `<processID>` with the process ID obtained in step 3:

```
sudo kill -9 <processID>
```
5. Run the following command:

```
sudo umount /data01/server_backups
```
6. For each Search Engine node, perform the following substeps:
 - a. Log in to the Search Engine node by using SSH.
 - b. Run the following command:

```
sudo umount /mnt/PPDM_Snapshots
```
7. Remove the DD system from PowerProtect Data Manager:
 - a. From the PowerProtect Data Manager UI, select **Infrastructure > Storage**.
 - b. Select the DD system to remove.
 - c. Click **Delete**.
8. Add the DD system back to PowerProtect Data Manager.
9. Enable server DR backups.

Troubleshooting NFS backup configuration issues

The following sections provide a list of error messages that might appear when you configure a server DR backup configuration that uses NFS.

DD storage unit mount command failed with error: 'Cannot mount *full path*: Access is denied'

This error message appears when an NFS export does not exist on the DD system for the full path to the server DR storage unit. This error message also appears when the redeployed virtual appliance was not added as a client for access to the NFS export.

To resolve this issue, ensure that you have configured an NFS export for the full path of the DD Boost storage unit and that the appliance is an Export client.

DD storage unit mount command failed with error: 'Cannot resolve *FQDN*: The name or service not known'

This error message appears when PowerProtect Data Manager cannot contact the DD system by using the specified FQDN. To resolve this issue, ensure that you can resolve the FQDN and IP address of the DD system.

Troubleshoot recovery of PowerProtect Data Manager

When the FQDN of the recovery site is different from the FQDN of the primary site, a mount error might occur and the recovery process requires a few extra steps.

About this task

If a mount error occurs during recovery, follow this work-around procedure.

Steps

1. On the DD system where the backup is located, delete the replication pair and mount it for PowerProtect Data Manager.
2. When recovery is complete, on PowerProtect Data Manager, regenerate the certificates using the following command.

```
sudo -H -u admin /usr/local/brs/puppet/scripts/generate_certificates.sh -c
```
3. Restart the system and select the URL of the primary PowerProtect Data Manager system.
The `https://PowerProtect Data Manager IP/#/progress` page appears and recovery resumes.
4. Log in to the primary PowerProtect Data Manager.
The PowerProtect Data Manager VM vCenter console shows an error, which you can ignore.
5. Open the primary PowerProtect Data Manager using the original IP address and log in.

Results

Recovery is complete.

Recover a failed PowerProtect Data Manager restore

Steps

1. Deploy a new PowerProtect Data Manager virtual appliance.
The *PowerProtect Data Manager Deployment Guide* for the appropriate platform provides instructions.
2. Contact Customer Support.

Disable server DR backups

Some maintenance procedures may require you to disable server DR backups during the procedure. Use this task only when referenced elsewhere.

Steps


1. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.
2. Click , select **Disaster Recovery**, and then click **Configuration**.
3. Record the existing server DR settings on the **Configuration** page.
4. Deselect **Enable backup**.
5. Click **Save**.

Next steps

After completing the maintenance procedure, re-enable server DR backups. The manual configuration procedure provides instructions.

Quick recovery for server DR

After a disaster, the quick recovery feature enables you to restore assets and data that you replicated to a destination system at a remote site.

 **NOTE:** Quick recovery does not re-create the original backup environment and source system which protected the restored assets. Therefore, quick recovery is not substitute for a server DR restore. To continue backing up the restored assets at the remote site, add the restored assets to a protection policy on the destination system.

Quick recovery is supported for the following assets protected by PowerProtect Data Manager:

- Virtual machines

 **NOTE:** This support does not include application-aware VADP workloads.

- Kubernetes namespaces and PVCs.
- File systems.

 **CAUTION:** Do not attempt the quick recovery of a system partition or boot disk.

Quick Recovery does not support restoring user data at the file or folder level.

Quick recovery sends metadata from the source system to the destination system, following the flow of backup copies. This metadata makes the replication destination aware of the copies and enables the recovery view. You can recover your workloads at the remote site before you have the opportunity to restore the source PowerProtect Data Manager system.

For example, the following figures show two sites that are named A and B, with independent PowerProtect Data Manager and DD systems for protection storage. Each site contains unique assets. Figure *Separate datacenters, before disaster* shows the initial configuration with both sites replicating copies to each other. Figure *Separate datacenters, after disaster* shows the aftermath, with site A down. The site A assets have been restored with quick recovery into the site B environment from the replicated copies.

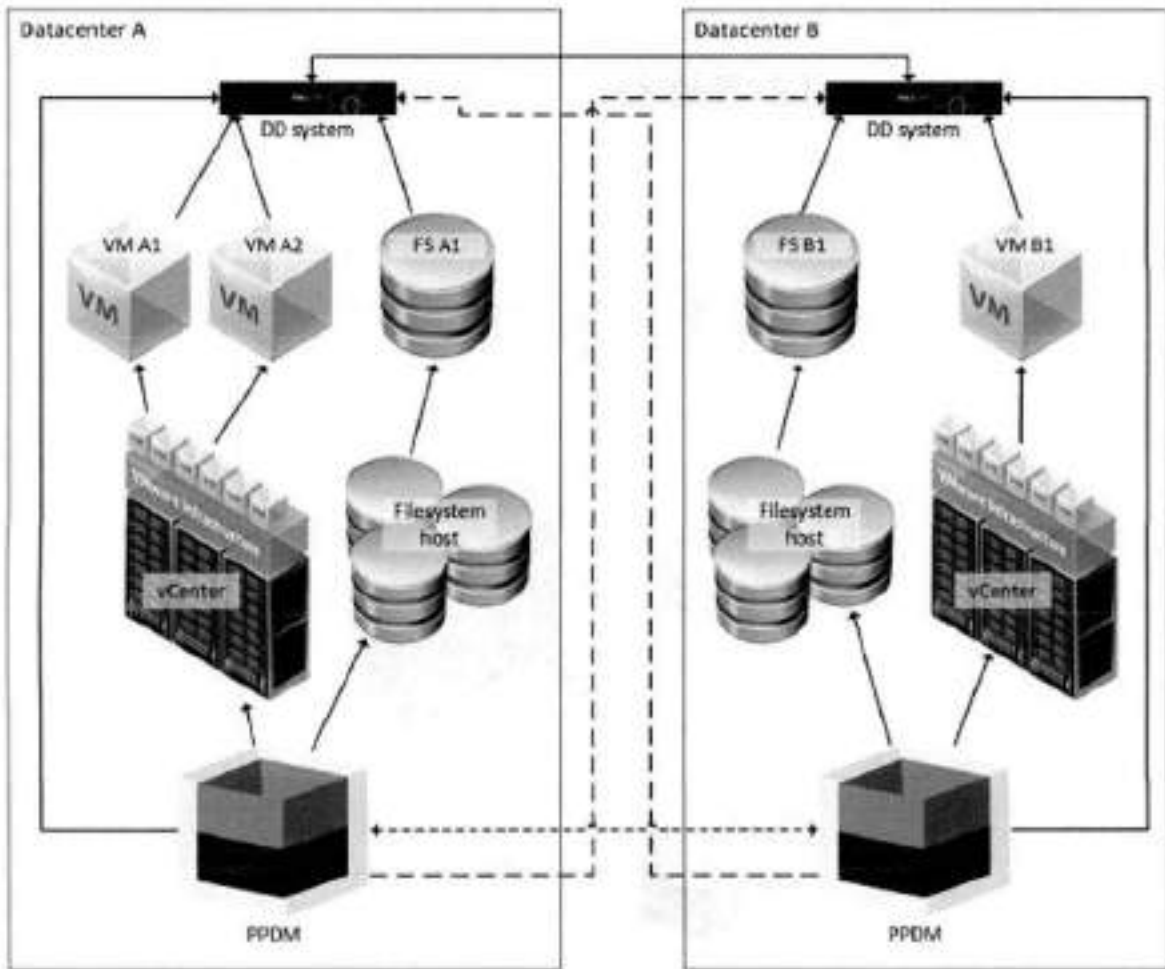


Figure 3. Separate datacenters, before disaster




Figure 4. Separate datacenters, after disaster

PowerProtect Data Manager supports quick recovery for alternate topologies. You can configure quick recovery for one-to-many and many-to-one replication. For example, the following figure shows a source PowerProtect Data Manager replicating to a standby DD system with its own PowerProtect Data Manager, all in the same data center. If the source system fails, the quick recovery feature ensures that you can still restore from those replicated copies before you restore the source.




Figure 5. Standby DD system

The following topics explain the prerequisites, how to configure PowerProtect Data Manager to support quick recovery, and how to use the recovery view to restore assets.

Quick recovery prerequisites

Before you configure quick recovery, complete the following items:

- Ensure that the source system and the destination system can ping each other using the same method: hostname or IP address.
- Ensure that the version of PowerProtect Data Manager is the same for both the source system and the remote (destination) system.
- Attach at least two protection storage systems to the source system; one for local protection storage and one for replication.
- Register asset sources with the source system and configure protection policies to protect those assets.
- Configure protection policies to replicate backup copies to the protection storage system at the remote site.
- Back up the protected assets and confirm that backup data successfully replicates to the destination protection storage system.
- Ensure that the replication protection storage is discovered in the remote (destination) system.
- Add and enable the asset source on the remote PowerProtect Data Manager instance.

- For Kubernetes quick recovery operations, ensure that the same Kubernetes cluster is not managed by more than one PowerProtect Data Manager instance.

Before you use the quick recovery remote view, add the destination system to the list of remote systems on the source. If you have changed the PowerProtect Data Manager security certificates after you configured quick recovery, the *PowerProtect Data Manager Security Configuration Guide* provides instructions to resynchronize the certificates.

Identifying a remote system

Remote systems added to PowerProtect Data Manager for quick recovery can be identified using either a fully qualified domain name (FQDN) or an Internet protocol (IP) address. If the incorrect identification is used, quick recovery fails with a certificate error.

If a remote system is already identified in the PowerProtect Data Manager certificate list, it must be added to PowerProtect Data Manager for quick recovery with the same identification.

If you always use either FQDNs or IP addresses for all remote systems, do the same for quick recovery.

If a certificate entry for the remote system exists, you must use the same identification when adding it for quick recovery. If you are unsure if a remote system you want to add for quick recovery is already in the PowerProtect Data Manager certificate list, perform the following steps:

- Log in to the console as the root user.
- Type `keytool -list -keystore`.
- Review the output and look for a certificate entry that corresponds to either the FQDN or IP address of the remote system.

Add a remote system for quick recovery

Configure PowerProtect Data Manager to send metadata to another system to which you have replicated backups. Only the Administrator role can add remote systems.

Steps

1. Click , select **Disaster Recovery**, and then click **Remote Systems**.
The **Remote Systems** tab opens and displays a table of configured remote PowerProtect Data Manager systems.
2. Click **Add**.
The **Add Remote PowerProtect System** window opens.
3. Complete the **Name** and **FQDN/IP** fields.
The **Name** field is a descriptive name to identify the remote system. To determine if you should enter the FQDN or IP address of the remote system, see *Identifying a remote system*.
4. In the **Port** field, type the port number for the REST API on the remote system.
The default port number for the REST API is 8443.
5. From the **Credentials** field, select an existing set of credentials with the Administrator role from the list.
Alternatively, you can click **Add Credentials** from this list to add new credentials with the Administrator role. Provide a descriptive name for the credentials, a username, and a password. Then, click **Save** to store the credentials.
6. Click **Verify**.
PowerProtect Data Manager contacts the remote system and obtains a security certificate for identity verification.
The **Verify Certificate** window opens to present the certificate details.
7. Review the certificate details and confirm each field against the expected value for the remote system. Then, click **Accept** to store the certificate.
The **Certificate** field changes to VERIFIED and lists the server's identify.
8. Click **Save**.
PowerProtect Data Manager returns to the **Remote Systems** tab of the **Disaster Recovery** window. The configuration change may take a moment to complete.
9. Click **Cancel**.
The **Disaster Recovery** window closes.
10. Click , select **Disaster Recovery**, and then click **Remote Systems**.

The **Remote Systems** tab opens.

11. Verify that the table of remote systems contains the new PowerProtect Data Manager system.
12. Click **Cancel**.
The **Disaster Recovery** window closes.

Next steps

On the remote system, enable the same asset sources that are enabled on this system. Enable an asset source provides more information. Enabling an asset source on the remote system makes replicated backups of that type visible and accessible.

On the remote system, open the recovery view and verify that backups are visible and accessible. It is recommended that you perform a test restore.

Metadata synchronizes between source and destination systems every three hours. If backups are not visible, allow sufficient time for the first synchronization before troubleshooting.

Edit a remote system

You can use the PowerProtect Data Manager user interface to change the descriptive name of the remote system, as well as the REST API port number and credentials. You can also enable or disable synchronization with the remote system. Only the Administrator role can edit remote systems.

Steps

1. Click  select **Disaster Recovery**, and then click **Remote Systems**.

The **Remote Systems** tab opens and displays a table of configured remote PowerProtect Data Manager systems.

2. Locate the row that corresponds to the appropriate remote system, and then select the checkbox for that row.
The PowerProtect Data Manager enables the **Edit** button.

3. Click **Edit**.

The **Edit Remote PowerProtect System** window opens.

4. Modify the appropriate parameters, and then click **Save**.

To enable or disable synchronization, select or deselect **Enable sync**. If you change the port number, you may need to re-verify the remote system security certificate.

PowerProtect Data Manager returns to the **Remote Systems** tab of the **Disaster Recovery** window. The configuration change may take a moment to complete.


5. Click **Cancel**.

The **Disaster Recovery** window closes.

Quick recovery remote view

Use the remote view to work with replicated copies on the destination system after the source is no longer available. For example, to restore critical assets before you are able to restore the source system.

On the destination system, log in as a user with the Administrator role. The remote server displays  on the banner.

When you click  and select **Remote Systems**, PowerProtect Data Manager presents a drop-down that contains the name of the local system and any connected systems. Each entry has the identifying suffix `(Local)` or `(Remote)`.

Select the source system from which you have replicated backups. PowerProtect Data Manager opens the remote view and presents a subset of the regular UI navigation tools:

- **Restore**
 - **Assets**— Shows replicated copies.
 - **Running Sessions**— Allows you to manage and monitor Instant Access sessions.
- **Alerts**— Shows alert information in a table, including audit logs.
- **Jobs**— Shows the status of any running restore jobs.

Each tool has the same function as for the local system. However, since the remote view is intended only for restore operations the scope is limited to the replicated copies from the selected source system. While in remote view, a banner identifies the selected system.

NOTE: For virtual machines, the quick recovery restore workflow does not include the **Restore VM Tags** option to restore vCenter tags and categories from the backup.

Use **Restore > Assets** to locate copies. The instructions for restoring each type of asset provide more information about restore operations.

When the recovery is complete, click **Remote Systems** and select the name of the local system to exit remote view.

Troubleshooting failed quick-recovery jobs

The following section provides troubleshooting information for when a quick-recovery job fails.

Quick-recovery jobs are not occurring

You might notice that expected quick-recovery jobs are not occurring.

If you review the `sync.log` file, you see entries similar to the following:

```
2022-07-14T00:10:03.235Z ERROR [] [scheduling-1] [][][][]
[c.e.b.s.e.r.i.ServerRestClient.checkHttpStatus(441)][441] - Return code = 401
UNAUTHORIZED, expected = 200 OK
2022-07-14T00:10:03.236Z ERROR [] [scheduling-1] [][][][]
[c.e.b.s.c.s.i.SyncHandshakeServiceImpl.syncHandshake(133)][133] - Failed to perform
handshake, version check failed.
com.emc.brs.sync.external.remote.VersionCheckResponseException: 500
INTERNAL_SERVER_ERROR "Unexpected error occurred during version check."; nested exception
is com.emc.brs.sync.external.remote.RemoteRestException: Incorrect credentials.
2022-07-14T00:10:03.236Z ERROR [] [scheduling-1] [][][][]
[c.e.b.s.c.s.i.SyncInstanceServiceImpl.syncHandshake(236)][236] - Failed to handshake
with the remote system: com.emc.brs.sync.model.SyncInstance
```

This error indicates that the password on the remote PowerProtect Data Manager system has changed.

To resolve this issue, change the credentials that are used for the remote PowerProtect Data Manager system.

Overview of PowerProtect Data Manager Cloud Disaster Recovery

The Cloud Disaster Recovery (DR) feature enables you to utilize a cloud DR site by deploying the Cloud DR Server in the public cloud. You can use the PowerProtect Data Manager UI for the purpose of running VM protection and DR workflows in the cloud.

Examples of Cloud DR workflows include the following:

- Cloud DR site copy management—Set the Cloud DR site by creating a VM protection policy in the PowerProtect Data Manager UI.
- VM copy failover validation—Before a disaster occurs, you can validate the failover of a VM copy to the cloud within PowerProtect Data Manager by running a DR test and then monitoring the test progress.
- Fail over a production VM—You can fail over a production virtual machine within PowerProtect Data Manager by running a DR failover operation and then verifying that the restored VM appears within Amazon Web Services (AWS) or Microsoft Azure cloud.
- Restore a production VM—You can restore virtual machines from copies that are stored in the cloud account (Amazon Web Services (AWS) or Microsoft Azure cloud) directly to vCenter. Restore operations are performed on one virtual machine at a time. You must manually select the target vCenter server.

The *PowerProtect Data Manager Cloud Disaster Recovery Administration and User Guide* provides more information about Cloud DR workflows within PowerProtect Data Manager.

Managing Alerts, Jobs, and Tasks


Topics:

- Configure Alert Notifications
- View and manage alerts
- View and manage Audit Logs
- Monitoring jobs and tasks
- Restart a job or task manually
- Restart a job or task automatically
- Resume misfire jobs after a PowerProtect Data Manager update
- Cancel a job or task
- Exporting logs
- Limitations for alerts, jobs, and tasks

Configure Alert Notifications

The **Alert Notifications** window of the PowerProtect Data Manager UI enables you to configure email notifications for PowerProtect Data Manager alerts.

Steps

1. From the PowerProtect Data Manager UI left navigation pane, select **Alerts**, and then select the **Alert Notifications** tab. The **Alert Notifications** window appears with a table that displays the details for existing notifications.
2. Click **Add**.
The **Add Alert Notification** dialog appears.
 **NOTE:** The **Add** button is disabled until you set up the email server. To add an alert notification, set up the email server in **System Settings > Support > Email Setup**. Set up the email server provides more information.
3. In the **Name** field, type name of the individual or group who will receive the notification email.
4. In the **Email** field:
 - a. Specify the email address or alias to receive notifications. This field is required in order to create an alert notification. Separate multiple entries with a comma.
 - b. Click **Test Email** to ensure that a valid SMTP configuration exists.
5. From the **Category** list, select one of the following notification categories:
 - All
 - Agent
 - Application Host Configuration
 - Cloud Tier
 - Compliance
 - Discover
 - Export Application Log
 - License
 - NAS Server Disaster Recovery
 - Protection
 - Protection Copy
 - Protection Infrastructure
 - Protection Policy
 - Protection Rule
 - Protection Source

- Push Update
 - Replication
 - Reporting
 - Restore
 - Security
 - Self Service
 - Server Disaster Recovery
 - System
6. From the **Severity** list, select one of the following notification severities:
 - All
 - Critical
 - Warning
 - Information
 7. In the **Duration** field, specify how often the notification email will be sent out. For example, you can set the duration to 60 minutes in order to send out a notification email every 60 minutes. If you set the duration to 0, PowerProtect Data Manager does not send out an email notification.
 8. In the **Subject** field, optionally type the subject that you would like to attach to the notification email.
 9. Click **Save** to save your changes and exit the dialog.

Results


The **Alert Notifications** window updates with the new alert notification. At any time, you can **Edit**, **Delete**, or **Disable** the notification by selecting the entry in the table and using the buttons in this window.

View and manage alerts

Alerts enable you to track the performance of data protection operations in PowerProtect Data Manager so that you can determine whether there is compliance to service level objectives. With the Administrator, Backup Administrator, Restore Administrator, or User role, you can access the alerts from the **Alerts** window. However, only some of these roles can manage alerts.

Steps



1. From the PowerProtect Data Manager UI left navigation pane, select **Alerts**.

You can also click  on the top banner, and then click the links to view unacknowledged alerts of all statuses (critical, warning, and informational), or only the unacknowledged critical alerts.


 **NOTE:** Clicking the **New** tag displays only the unacknowledged alerts that have been generated within the last 24 hours.

The number that appears next to  is the total number of unacknowledged critical alerts over the last 24 hours.

The **Alerts** window displays.

2. Select the **System** tab. A table with an entry for each applicable alert displays. By default, only unacknowledged critical alerts from the last 24 hours display, unless you selected to view all unacknowledged alerts from the links under . If filter tags have already been applied, the window displays these filter tags. Click **X** next to any of these filter tags to clear a filter, and the table view updates with the applicable selections. You can sort the alerts in the table by Severity (Critical, Warning, Informational), Date, Category, or Status (Acknowledged or Unacknowledged).
3. Select a time from the last 24 hours, the last 3 days, the last 7 days, the last 30 days, and a specific date for the alerts you want to view, or provide a custom time range. You can also select **All Alerts** from this list to display information for all alerts that match the filter tags.
4. Optionally, clear the **Show only unacknowledged alerts** checkbox if you want to view both acknowledged and unacknowledged alerts. If you clear this checkbox, the **Unacknowledged** filter tag is also cleared.
5. To view more details about a specific entry, click  next to the entry in the table.
6. For the following steps, log in to the PowerProtect Data Manager UI with an account that has the Administrator, Backup Administrator, or Restore Administrator role.

7. To acknowledge one or more alerts, select the alerts and then click **Acknowledge**.
8. To add or edit a note for the alert, click **Add/Edit Note**, and when finished, click **Save**.
9. To export a report of alert information to a `.csv` file which you can download for Excel, click **Export All**.

 **NOTE:** If you apply any filters in the table, exported alerts include only those alerts that satisfy the filter conditions.

View and manage Audit Logs


Audit logs enable you to view specific information about jobs that are initiated in PowerProtect Data Manager so that you can determine compliance to service level objectives. You can access the audit logs from the **Administration > Audit Logs** window.

Steps

1. From the PowerProtect Data Manager UI left navigation pane, select **Administration > Audit Logs**.

The **Audit Logs** window displays audit information in a table.

2. (Optional) Sort and filter audit information:

- To filter audits by **Audit Type**, **Changed By**, or **Object Changed**, click .
- To sort audits by **Changed At**, **Audit Type**, **Changed By**, or **Object Changed**, click a column heading.
- To filter audits based on a search string, type a keyword in the **Search** field.

3. To view more details about a specific entry, click  next to the entry in the table.

- Review the information for the audit log.
- Optionally, add a note for this audit log in the **Notes** field.

4. To export an audit log report to a `.csv` file which you can download as an Excel file, click **Export All**.

 **NOTE:** If you apply any filters in the table, exported audit logs include only those logs that satisfy the filter conditions.

5. To change the retention period for audit logs, click **Set Boundaries**, select the number of days from the **Days of Retention** menu, and then click **Save**.


Monitoring jobs and tasks

For jobs, the PowerProtect Data Manager UI provides three window views based on the job type — **Protection Jobs**, **Asset Jobs**, and **System Jobs**. These windows allow you to monitor the status of data protection, system, and maintenance jobs, and view details about failed, in progress, or recently completed jobs. To perform analysis or troubleshooting, you can view a detailed log of a failed job or task, and also see any errors that occurred.

Use the filtering and sorting options in each window to find specific jobs or tasks, and to organize the information that you see. Filter, group, and sort jobs provides more information. You can also view details for a job group in addition to individual jobs and tasks:

- For protection and system jobs, when you click the job ID next to the job entry, the **Job ID Summary** window displays the information for only this job group, job, or task.
- For asset jobs, when you select the row for the job in the table, a pane opens at the right of the window that displays information for this asset job.

From these views, you can monitor the status of individual jobs and tasks in the **Step Log** tab, view job and task details in the **Details** tab, and, if eligible, perform certain operations on jobs and tasks such as **Restart** or **Cancel**.

 **NOTE:** The **Jobs** windows have been optimized for a screen resolution of at least 1920 x 1080, with 100% scaling. Display issues might occur for smaller screens. Set your screen resolution to at least 1920 x 1080, with 100% scaling.

Monitor and view jobs and assets

Use the **Protection Jobs**, **Asset Jobs** and **System Jobs** windows to monitor and view status information for PowerProtect Data Manager operations.

Within these windows, you can export job records and asset activities by using the **Export All** functionality.

Protection jobs

To view protection jobs and job groups, from the PowerProtect Data Manager UI left navigation pane, select **Jobs > Protection Jobs**.

The **Protection Jobs** window opens to display a list of protection jobs and job groups.

Protection jobs include:

- Cloud Tier
- Cloud Protect
- Consolidated Cloud Snapshot Manager jobs
- **i** **NOTE:** This job type does not apply to SAP HANA databases.
- Export Reuse
- Indexing
- Protect
- Replicate
- Restore

You can monitor and view detailed information for both centralized and self-service backup and restores of database application assets.

i **NOTE:** The **Cancel** and **Retry** options are not available for self-service jobs that are created by database application agents.

For application assets, the **Protect**, **Restore**, and **Replicate** job types can be monitored at the host or individual asset level. For all other asset types, the **Protect** and **Replicate** job types can be monitored at the host or individual asset level.

Asset jobs

The **Asset Jobs** window allows you to view all jobs for a specific asset or application agent host, and to view the history of protection activities at the asset/agent host level.

To view information about assets for which jobs have been run, from the PowerProtect Data Manager UI left navigation pane, select **Jobs > Asset Jobs**.

The **Asset Jobs** window opens to display a list of assets. For application agent assets, you can also view the associated host. You can filter by asset/host name or by job type.

Examples of asset job types include:

- Application Host Configuration
- Cloud Copy Recover
- Cloud Disaster Recovery
- Cloud Protect
- Cloud Tier
- Config
- Delete
- Disaster Recovery
- Export Reuse
- Indexing
- Manage
- Notify
- Protect
- Push Update
- Replicate

- Restore
- System
- Validate

NOTE: The PowerProtect Data Manager UI **Dashboard** additionally provides details for any successful, partially successful, failed and canceled jobs at the asset/host level.

System jobs

To view system jobs and job groups, from the PowerProtect Data Manager UI left navigation pane, select **Jobs > System Jobs**.

The **System Jobs** window opens to display a list of system jobs and job groups.

System jobs include:

- Config
- Console
- Delete
- Disaster Recovery
- Cloud Disaster Recovery
- Cloud Copy Recovery
- Discovery
- Manage
- Notify
- System
- Validate

System jobs can be monitored at the job group or job level.

Job information



The main **Protection Jobs** and **System Jobs** windows lists basic job information.

The following information is available in the **Protection Jobs** and **System Jobs** windows.

Table 30. Job information

Column	Description
Job ID	The unique and searchable identifier for the job.
Status	Indicates the current state of the job. A job can be in one of the following states: <ul style="list-style-type: none"> • Success • Completed with Exceptions • Failed • Canceled • Unknown • Skipped • Running • Queued • Canceling For jobs that do not have a Success status, a count of jobs is shown next to the status.
Description	Description of the job.
Policy Name	Name of the protection policy that started the job.
Assets	Number of individual assets or tasks within the job group.
Job Type	Type of protection job or system job.
Asset Type	Type of asset.
Start Time	Date and time that the job is scheduled to begin.

Table 30. Job information (continued)

Column	Description
End Time	Date and time that this job completed. This column is not shown by default. To see a complete list of filtering and sorting columns, click  .
Duration	Overall duration of the job. This column is not shown by default. To see a complete list of filtering and sorting columns, click  .

View details for protection jobs

In the **Job ID Summary** window for protection jobs, you can view details and status of specific jobs. For application protection jobs, you can view details and status of specific jobs and assets. This information can be helpful when troubleshooting to determine whether one or more assets caused a job to fail.

Steps

1. From the PowerProtect Data Manager UI left navigation pane, select **Jobs > Protection Jobs**.
2. Click the job ID next to the job name.

The **Job ID Summary** window opens and lists all jobs as entries in the table.

You can filter, group, and sort the information that appears in the window. Filter, group, and sort jobs provides more information.

The policy name, job type, and asset type appear at the top of the **Job ID Summary** window.

The overall job group metrics and details also appear, as shown in the following figure.



Figure 6. Job Metrics and Job Details

The **Job Metrics** section displays the number of assets, the total size of the data transferred, and the overall duration of the job group. The total duration of jobs within the job group is shorter than the duration indicated in the job metrics. When you restart a protection job that is part of a completed job group, the duration that is indicated in the job metrics does not include the time that is elapsed between when the job group completed and when the job was restarted. In addition, it does not include the time that it takes for the retried job to run.

The **Job Details** section displays more specific information such as the job start and end time, the protection storage target, the average data transfer rate, the amount of data changed since the last protection job, the average throughput, and the rate of compression applied. For restore jobs of Microsoft SQL Server databases, some fields are either not applicable or set to zero.

Job metrics and details do not display or might be incomplete for job groups that contain Oracle database assets.

Click **Hide Summary** to hide job metrics and details, or click **Show Summary** to view job metrics and details.


When you hover over a job, the **Job ID Summary** displays a message for the job to indicate its progress. Depending on the job and if any issues are detected, one of the following statuses is shown:



- **No reported issues**—No issues affecting the job.
- **Timeout issues**—Timeout issues might be affecting the job.
- **Connectivity issues**—Network connectivity issues might be affecting the job.
- **State stall issues**—Progress for this job is stalled.

The **Job ID Summary** window provides summary data for specific jobs and assets in a table view. For grouped assets, the host-level entry indicates the sum of the values of a given metric for every asset on the host.

The following table describes the columns that might appear in the window. Not all columns appear in the **Job ID Summary** window of every asset type.

Table 31. Job ID Summary window details

Column	Description
Details	Click  in the Details column to view job statistics and summary information.
Asset	Name of the job for the asset.
Status	Indicates the current state of the job. A job can be in one of the following states: <ul style="list-style-type: none"> • Success • Completed with Exceptions • Failed • Canceled • Unknown • Skipped • Running • Queued • Canceling
Size	Size of job for the asset.
Data Transferred	Total data that is transferred to storage.
Reduction %	Total reduction percentage of storage capacity for the job.
Start Time	Date and time that the job is scheduled to begin.
End Time	Date and time that this job completed.
Error Code	If the job did not successfully complete, a numeric error code appears. To view a detailed explanation, double-click the error code.
Host/Cluster/Group Name	The hostname, cluster, or group name that is associated with the asset.
Duration	Overall duration of the job. This column only appears for Protect and Replicate job types for application assets.
Asset Size	Total size of the asset in bytes.
Data Compressed	Capacity that is used after client compression of the data in bytes. This column only appears for Protect and Replicate job types for application assets.
Download log	Detailed log for an asset or task that you can export and download.

3. To view job details and summary information, click  in the **Details** column next to the job, or expand the entry for the job group by clicking .

For grouped assets, the **Job ID Summary** window lists the individual jobs for each asset within the job group.

The right pane appears and displays the following information about the job or task:

- **Step Log**—Displays a list of steps that have been completed or are in progress for the job or task, and indicates the amount of time that was required to complete each step. If a job step is still active, the Step Log also provides a more detailed description about what aspect of the step is being performed.
- **Details**—Displays statistics and summary information, such as the start time and end time, asset size, duration, and additional details.
- **Error**—Displays error details for failed jobs.
- **Canceled**—Displays details for canceled jobs.
- **Skipped**—Displays details for skipped jobs.
- **Unknown**—Displays details for jobs with an unknown status.

View details for asset jobs

In the right pane of the **Asset Jobs** window, you can view details and status information for assets that have been included in active, completed or failed PowerProtect Data Manager jobs. This information can be helpful when tracking the progress of a job, or when troubleshooting to determine why the configuration or protection of a particular asset was unsuccessful.

About this task

If a job is in progress or has been performed for an asset within the last 45 days, the asset appears with a link in the **Infrastructure > Assets** window. When clicked, this link opens the **Jobs > Asset Jobs** window.

Steps

1. From the PowerProtect Data Manager UI left navigation pane, select **Jobs > Asset Jobs**.
By default, the table displays a list of assets for which jobs have been run in the last 24 hours.

The following table describes the asset job details that might appear, depending on which columns have been customized.

Table 32. Asset Jobs window details

Column	Description
Asset	Name of the asset in the protection job.
Host	For application agent assets, the hostname that is associated with the asset.
Status	Indicates the current state of the job. A job can be in one of the following states: <ul style="list-style-type: none">• Success• Completed with Exceptions• Failed• Canceled• Unknown• Skipped• % (indicating the progress of the job)• Queued• Canceling
Policy Name	The protection policy that contains this asset.
Job Type	Supported asset job types include Config, Protect, Replicate, Restore, and Cloud Tier
Asset Type	Indicates the specific type of asset. For example, VMware Virtual Machine.
Start Time	Date and time that the job is scheduled to begin.
Duration	Overall duration of the job.
Details	Select the row of the asset to open the Details tab in the right pane, where you can view statistics and summary information.


Table 32. Asset Jobs window details (continued)


Column	Description
Step Log	From the right pane, select the Step Log tab to view a list of steps that ave been completed for the asset job, along with the amount of time that was required to complete each step.
Errors	If the job did not complete successfully, select the row of the asset to open the Errors tab in the right pane, where you can view any errors along with a numeric error code.

2. Optionally, customize the asset jobs that display:
 - a. Select a different time period or specify a time range by clicking the **Start Time** box.
 - b. Use the filter in each column to display only assets that match the search criteria.
 - c. Click a status in the window's summary information to view only assets with a particular job status.
 - d. Sort the information by clicking the up and down arrows within each column.

When the view is customized, the time range, search filter and status filter persist in the PowerProtect Data Manager UI until the filters are cleared. Filter, group, and sort jobs provides more information.

3. Select the row of the asset job.

A pane displays to the right of the window, as shown in the following figure. At any time, click  at the top of the pane to hide or show the details. This pane displays the following tabs:

- **Step Log**—Displays a list of steps that have been completed or are in progress for the asset job, and indicates the amount of time that was required to complete each step. If a job step is still active, the Step Log also provides a more detailed description about what aspect of the step is being performed.
 -  **NOTE:** The Step Log and description only displays for jobs related to backup, restore, and disaster recovery operations.
- **Details**—Displays statistics and summary information, such as the start time and end time, asset size, duration, and additional details.
- **Error**—Displays any errors that occurred if the asset job failed or completed with exceptions

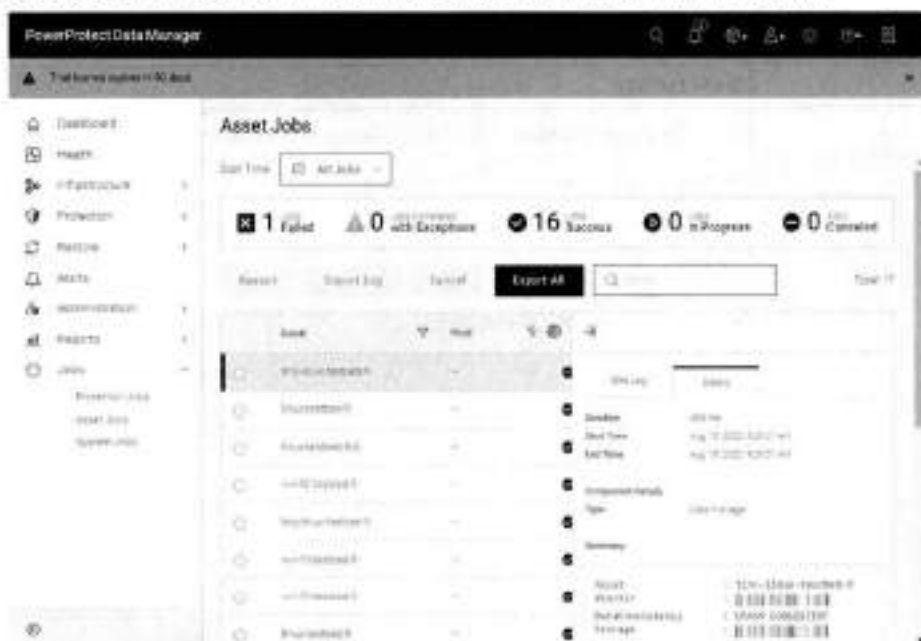


Figure 7. Asset details, step log, and errors

4. If the job failed, was canceled, or completed with exceptions, and is eligible for restarting, select the radio button next to the asset and click **Restart**.
5. To export the step log for an asset job, select the radio button next to the asset job and click **Export Log**, or click **Export All** to create a .csv file for all asset jobs.

View details for system jobs and tasks

In the **Job ID Summary** window for system jobs, you can view details and status of specific jobs and tasks. This information can be helpful when troubleshooting to determine whether one or more jobs or tasks caused a job to fail.

Steps

1. From the PowerProtect Data Manager UI left navigation pane, select **Jobs > System Jobs**.
2. Click the job ID next to the job name.

The **Job ID Summary** window opens to display a list of all system jobs or tasks.

You can filter, group, and sort the information that appears in the window. Filter, group, and sort jobs provides more information.

For jobs and tasks, a table appears at the bottom of the window. The success or failure of individual tasks is indicated in the **Status** column. If a failed job or task requires action, a status of **Critical** appears.


You can view job status and summary information for scheduled discovery of application assets and application systems. If a discovery job fails, PowerProtect Data Manager displays error details and steps to resolve the issue. An alert is also generated in the **Alerts** window.

When you hover over a job or task, the **Job ID Summary** displays a message for the job to indicate its progress. Depending on the job and if any issues are detected, one of the following statuses is shown:

- **No reported issues**—No issues affecting the job.
- **Timeout issues**—Timeout issues might be affecting the job.
- **Connectivity issues**—Network connectivity issues might be affecting the job.
- **Stats stall issues**—Progress for this job is stalled.

The **Job ID Summary** window provides summary data for specific jobs and tasks in a table view. The following table describes the columns that might appear in the window. Not all columns will appear in the **Job ID Summary** window of every asset type.

Table 33. Job ID Summary window details

Column	Description
Details	Click  in the Details column to view job or task statistics and summary information.
Task Name	Name of the task.
Status	Indicates the current state of the job or task. A job or task can be in one of the following states: <ul style="list-style-type: none">• Success• Completed with Exceptions• Failed• Canceled• Unknown• Skipped• Running• Queued• Canceling
Asset	Name of the asset.
Start Time	Date and time that the job or task is scheduled to begin.
Duration	Overall duration of the job or task.
Data Transferred	Total data that is transferred to storage.

3. To view job or task details and summary information, click  in the **Details** column next to the individual job or task.

The right pane appears and displays the following information about the job or task:

- **Step Log**—Displays a list of steps that have been completed or are in progress for the job or task, and indicates the amount of time that was required to complete each step. If a job step is still active, the Step Log also provides a more detailed description about what aspect of the step is being performed.

NOTE: The Step Log and description only displays for jobs related to backup, restore, and disaster recovery operations.

- **Details**—Displays statistics and summary information, such as the start time and end time, asset size, duration, and additional details.
- **Error**—Displays error details for failed jobs.
- **Canceled**—Displays details for canceled jobs.
- **Skipped**—Displays details for skipped jobs.
- **Unknown**—Displays details for jobs with an unknown status.

Filter, group, and sort jobs

The **Protection Jobs**, **Asset Jobs**, and **System Jobs** windows provide options to filter, group, and sort the information that appears.

Filter jobs by status

Use the quick filters at the top of the window to filter jobs by status. By default, all jobs are shown regardless of status. To display only jobs with a specific status, at the top of the window, select one of the following options:

- **Failed**
- **Completed with Exceptions**
- **Success**
- **Canceled**
- **In Progress**

In Progress jobs include **Running**, **Queued**, and **Canceling** jobs.

When you select a quick filter to filter jobs by a certain status, the window displays the filter above the table. To stop filtering by the selected status, click **x**.

Filter jobs by start time

Use the **Start Time** filter to display jobs that started in a specified period. Jobs are retained for a maximum of 45 days. Select from one of the following options:

- All jobs
- Last 24 hours
- Last 3 days
- Last 7 days
- Last 30 days
- Specific date
- Custom date range

Group jobs

In the **Protection Jobs** and **System Jobs** windows, select a job to display its **Job ID Summary** window. The **Group by** feature in the **Job ID Summary** window provides options to group assets within a protection job.

The following asset types support the **Group by** feature:

- Microsoft SQL Server databases
- Microsoft Exchange Server databases
- Oracle databases
- File Systems
- SAP HANA databases
- Kubernetes clusters

- Network-attached storage (NAS) shares
- VMware Virtual Machines

To group assets in a protection job, in the **Job ID Summary** window for the job, select an option from the **Group By** drop-down list. To display all assets, select **Group by > None**. For example, to group virtual machine assets by ESX host, click **Group by > ESX Host**.

The following table lists the available **Group by** options:

Table 34. Group by options


Asset type	Options
Microsoft SQL Server database	SQL Host
	SQL Instance
Oracle database	Oracle Host
	Oracle Instance
File System	File System Host
	File System Host OS
Microsoft Exchange Server database	Exchange Host
SAP HANA database	SAP HANA Host
Kubernetes	Kubernetes Cluster
	Kubernetes Namespace
NAS	NAS Server
	NAS Appliance
VMware Virtual Machine	Datastore
	ESX Host
	Virtual Datacenter
	VM Guest OS
	VMware Cluster

NOTE: Currently, the **Group by** filter is only available for the **Protect** job types.

Search filter

Use the **Search** field to filter jobs based on a search string. When you type a keyword in the **Search** field, the PowerProtect Data Manager UI filters the results as you type. To clear the search filter, remove all keywords from the **Search** field.

Filter and sort information in tables

You can filter and sort the information that appears in table columns. Click  in the column heading to filter the information in a table column, or click a table column heading to sort that column.

To see a complete list of filtering and sorting columns, click . Depending on the type of job, the available filtering and sorting columns might differ.

The following filtering and sorting options are available for jobs and tasks:

Table 35. Protection, Asset ,and System Jobs windows

Filtering options	Sorting options
Filter jobs or tasks by Job ID, Status, Description, Policy Name, Job Type, End Time, and Asset Type .	Sort jobs or tasks by Job ID, Description, Policy Name, Job Type, Asset Type, Start Time, and End Time .

Table 36. Job ID Summary window for protection jobs

Filtering options	Sorting options
Filter jobs by Asset, Status, Error Code, Start Time, or End Time . For application assets, you can also filter jobs by Host/Cluster/Group Name . ⓘ NOTE: For application assets, these options are only available when you select Group by > None .	Sort jobs by Asset, Status, Error Code, Size, Data Transferred, Reduction %, Start Time, End Time, or Duration . For application assets, you can also sort jobs by Host/Cluster/Group Name . ⓘ NOTE: For application assets, these options are only available when you select Group by > None .

Table 37. Job ID Summary window for system jobs

Filtering options	Sorting options
Filter jobs or tasks by Task Name, Status, Asset, or Start Time .	Sort jobs or tasks by Task Name, Status, Asset, Start Time, Duration, or Data Transferred .

Restart a job or task manually

You can manually restart a failed virtual machine backup.

About this task

When you click **Restart**, the job or task restarts immediately, regardless of the scheduled activity window.

Note the following:

- If a policy with both protection and Cloud Data Recovery objectives fails, the Cloud Data Recovery job is canceled and cannot be restarted.
- Cloud Snapshot Manager jobs cannot be restarted.

Steps

1. From the PowerProtect Data Manager UI left navigation pane, select **Jobs > Protection Jobs**, **Jobs > Asset Jobs**, or **Jobs > System Jobs**.
The window displays all completed and running jobs.
2. To restart a failed job or job group, select the failed job or job group from the list, and then click **Restart**. If the job is ineligible for restart, the button will be grayed out.
3. To restart a failed system or protection job or task from the **Job ID Summary** window:
 - a. Click the job ID next to the name of the job or job group.
The **Job ID Summary** window opens to display a list of all jobs or tasks.
 - b. Select the job or task from the list, and then click **Restart**.

Results

After the job or task has been restarted, the status indicates **Running** or **Queued**.

- ⓘ **NOTE:** When you restart a protection job that is part of a completed job group, the duration indicated in the **Job Metrics** includes the time that elapsed between when the job group completed and when the job was restarted, in addition to the time it takes for the retried job to run.

Restart a job or task automatically

If a backup job fails or one of the tasks within the job fails, you can enable automatic restart of the failure by configuring auto retry in the `entrypoint.sh` file. Auto retry can be useful in situations where the failure is due to an intermittent issue, such as a network or service interruption.

Prerequisites

In PowerProtect Data Manager, some services that are required for auto retry, such as the workflow service, have been moved into a docker container. In order to enable auto retry, ensure that the workflow service is running in a docker.

About this task

Auto retry is only supported for daily, weekly, or monthly schedules for virtual machine and File System agent protection operations.

Steps

1. Log in to the PowerProtect Data Manager server by using SSH.
2. Copy the `entrypoint.sh` file from the workflow container by typing the following:

```
docker cp workflow:/workflow/bin/entrypoint.sh .
```
3. Configure auto retry by adding a line to `entrypoint.sh`:
 - a. Type `vi entrypoint.sh`
 - b. Before the last line in the output, add the following:

```
-Denable.auto.retry.scheduler=true \
```

NOTE: Auto retry is disabled by default. After adding this line, if you want to disable this setting at any point, change the entry to `-Denable.auto.retry.scheduler=false \`
4. Optionally, add the following application properties to the file to specify a maximum number of auto retries and a time interval at which subsequent auto retry attempts will occur:

```
-Dfailed.job.retry.max.count=2 \  
-Dfailed.job.retry.interval=PT30M \
```

NOTE: The values specified above are the recommended default values. Auto retries will only occur during the activity window. If you perform a manual retry in the PowerProtect Data Manager UI, this retry will not count towards the auto retry max count.

For the interval duration, the value must be specified in ISO-8601 format.
5. Save the `entrypoint.sh` file to the workflow container by typing the following:

```
docker cp entrypoint.sh workflow:/workflow/bin/
```
6. Restart the workflow service by using one of the following methods:
 - Type `docker container restart workflow`

NOTE: For the configuration to be applied successfully using this method, you can only restart the container. If you restart your workflow service or your PowerProtect Data Manager operating system, the configuration will be lost.
 - Type the following to save the docker image and restart the workflow service. For example:

```
docker commit workflow dpd/ppdm/ppdm-workflow:PowerProtect Data Manager version  
workflow restart
```

where *PowerProtect Data Manager version* is the PowerProtect Data Manager version that is deployed on your system.
You can use this method to permanently apply the configuration change after restoring the docker image.

Results

After configuration, the workflow service is scheduled to run every 30 minutes to determine if any jobs or tasks have failed. If a restart occurs, the status indicates **Running** or **Queued**. To view whether a failed job or task has restarted, go to the **Jobs** window in the PowerProtect Data Manager UI and select **Running** or **Queued**.

Resume misfire jobs after a PowerProtect Data Manager update

During an update, the PowerProtect Data Manager system enters maintenance mode. Any job that is not in queue and is scheduled to run during the time that the PowerProtect Data Manager system is in maintenance mode will be missed. These missed jobs are known as misfires. As of this release, PowerProtect Data Manager uses the Quartz Scheduler to resume scheduled workflows when the service recovers or when the schedule resumes.

About this task

The trigger and firing data of jobs are stored in a database application. If the schedule service is down, such as during an update, the Quartz Scheduler recovers this data and resumes the jobs when the PowerProtect Data Manager system is operational again.

NOTE: In the current release, this feature is enabled by default.

You can enable or disable the misfire feature by configuring the `entrypoint.sh` file.

Steps

1. Log in to the PowerProtect Data Manager server by using SSH.
2. Copy the `entrypoint.sh` file from the scheduler container by typing the following:

```
docker cp scheduler:/scheduler/bin/entrypoint.sh .
```

3. Configure the misfire conditions in the `entrypoint.sh` file:

NOTE: Before the last line in the output, `-jar /${APP_NAME}/lib/scheduler-core.jar`, add the lines for each misfire condition.

- a. To enable misfire and trigger each job once, add the following properties and corresponding values:

```
-Dspring.quartz.properties.misfire.cron.strategy=WITH_MISFIRE_HANDLING_INSTRUCTION_FIRE_AND_PROCEED \
```

NOTE: This condition is enabled by default.

```
-Dspring.quartz.properties.misfire.calendar.strategy=WITH_MISFIRE_HANDLING_INSTRUCTION_FIRE_AND_PROCEED \
```

- b. To enable misfire and trigger each job as many times as misfire happens, add the following properties and corresponding values:

```
-Dspring.quartz.properties.misfire.cron.strategy=WITH_MISFIRE_HANDLING_INSTRUCTION_IGNORE_MISFIRES \
```

```
-Dspring.quartz.properties.misfire.calendar.strategy=WITH_MISFIRE_HANDLING_INSTRUCTION_IGNORE_MISFIRES \
```

- c. To disable misfire, add the following properties and corresponding values:

```
-Dspring.quartz.properties.misfire.cron.strategy=WITH_MISFIRE_HANDLING_INSTRUCTION_DO_NOTHING \
```

```
-Dspring.quartz.properties.misfire.calendar.strategy=WITH_MISFIRE_HANDLING_INSTRUCTION_DO_NOTHING \
```

4. Save the `entrypoint.sh` file to the scheduler container by typing the following:

```
docker cp entrypoint.sh scheduler:/scheduler/bin/
```

5. Restart the scheduler service by using one of the following methods:

- Type `docker container restart scheduler`

NOTE: For the configuration to be applied successfully using this method, you can only restart the container. If you restart your scheduler service or your PowerProtect Data Manager operating system, the configuration will be lost.

- Type the following to save the docker image and restart the scheduler service:

```
docker commit scheduler dpd/ppdm/ppdm-scheduler:PowerProtect Data Manager version
```

scheduler restart

where *PowerProtect Data Manager version* is the PowerProtect Data Manager version that is deployed on your system. You can use this method to permanently apply the configuration change after restoring the docker image.

i **NOTE:** Ensure that the PowerProtect Data Manager version specified in the `commit` command matches the PowerProtect Data Manager version that is deployed on your system.

Cancel a job or task

From the PowerProtect Data Manager UI, you can cancel a backup or restore that is still in progress, or any asset protection and replication activities when the tasks are queued.

About this task

i **NOTE:** The **Cancel** operation is available for the following supported jobs and tasks only:

- Backup and restore of:
 - Virtual machine assets
 - Kubernetes assets
 - NAS assets
 - File System assets
 - Microsoft SQL Server assets
 - Block volume assets
 - Server DR
 - Cloud DR
- Backup (only) of:
 - Microsoft Exchange Server assets
 - Oracle assets
 - SAP HANA assets
 - Transaction logs of application-aware asset backups
- Replication
- Compliance
 - Copy deletion
 - Compliance verification
 - Auto promotion to full backup
 - Cleaning MTree or deleting user
 - On-demand update retention
- Support
 - Communication of telemetry data
 - Export of job and job group logs
 - Adding log bundles

Steps

1. From the PowerProtect Data Manager UI left navigation pane, select **Jobs > Protection Jobs**, **Jobs > Asset Jobs**, or **Jobs > System Jobs**.

The relevant **Jobs** window appears, displaying all completed and running jobs.

2. To cancel a job or job group, select a job or job group that is in-progress, and then click **Cancel**.

i **NOTE:** If a job is almost complete, the cancellation might fail. If the cancellation fails, a message displays indicating that the job cannot be canceled.


The window displays the status of the canceled job or job group. If the cancellation is successful, then the status eventually changes to **Canceled**. If the cancellation is not successful, then the status might indicate either **Success** or **Critical**.

3. For protection and system jobs, to cancel an individual job or task from the **Job ID Summary** window:

- a. Click the job ID next to the name of the job or job group.

The **Job ID Summary** window opens to display a list of all jobs or tasks.

- b. Select a job or task that is in-progress, and then click **Cancel**.


 **NOTE:** If a job or task is almost complete, the cancellation might fail. If the cancellation fails, a message displays indicating that the task cannot be canceled.

- c. Click **Close**.

The **Job ID Summary** window displays the status of the canceled job or task. If the cancellation is successful, then the status eventually changes to **Canceled**. If the cancellation is not successful, then the status might indicate either **Success** or **Critical**.

Exporting logs

The PowerProtect Data Manager UI enables you to export and download a detailed log of a job, asset, or task to perform analysis or troubleshooting.

You can export and download a log for a job, asset, or task with any status. After you export a log, you can download it by clicking .

Export logs for jobs

You can export and download a log for a protection job or system job by using the PowerProtect Data Manager UI.

About this task

PowerProtect Data Manager restricts the log export function in the following situations:

- The job is from a different PowerProtect Data Manager tenant.
- The job supports exporting an external log at the current stage for the following asset sources:
 - Virtual machines
 - Kubernetes
 - Microsoft SQL Server
 - Microsoft Exchange Server
 - File Systems
 - Oracle
 - SAP HANA
 - Network-attached storage (NAS)

In these situations, create a log bundle instead. In the PowerProtect Data Manager UI, select **Settings > Support > Logs** to add a log bundle.

Steps

1. From the PowerProtect Data Manager UI left navigation pane, select **Jobs > Protection Jobs**, **Jobs > Asset Jobs**, or **Jobs > System Jobs**.

The relevant **Jobs** window appears, displaying all jobs.

2. Select a job from the list, and then click **Export Log**.

Hover over  next to the asset or task in the **Download Log** column to display the progress. When the log export is complete, you can download the log.

3. Click  next to the ID for the job to download the exported log.

Export logs for assets or tasks

You can export and download a log for an individual asset or task.

Steps

1. From the PowerProtect Data Manager UI left navigation pane, select **Jobs > Asset Jobs**.
The **Asset Jobs** window appears.
2. Select the row of the asset, and then click **Export Log**.

Hover over  next to the asset or task in the **Download Log** column to display the progress. When the log export is complete, you can download the log.

3. Click  in the **Download Log** column to download the exported log.

Limitations for alerts, jobs, and tasks

Review the following limitations that are related to alerts, jobs, and tasks.

For in-progress jobs, the details pane displays the "Error" tab and indicates "Failed"

When you open the **Details** pane for in-progress jobs, the **Error** tab appears and incorrectly indicates **Failed** in the error details.

Workaround

Ignore the **Error** tab for in-progress jobs.

Self-service jobs are not showing on PowerProtect Data Manager Protection Jobs window after recreating lockbox entry

On both Windows and Linux, self-service jobs do not appear on PowerProtect Data Manager Protection Jobs window after recreating a lockbox entry.

Workaround

Restart the agent service or change the system time (+24 hrs).

The history of viewable backup jobs is limited to the 10,000 most recent

If viewing a history of backup jobs and trying to move to a page in the interface that would show the 10,000th or earlier backup job, the following error is seen:

```
error: 416: "The query will return too many results."
```

Workaround

To view earlier backup jobs, use a filter that includes the earlier backup jobs but limits the number of entries to less than 10,000.

Total protection jobs count on the PowerProtect Data Manager dashboard does not include skipped jobs

The **Total Jobs** count shown in the **Jobs | Protection** widget on the dashboard does not include skipped jobs. As a result, this count does not reflect the total count of protection jobs that is shown in the **Protection Jobs** window.

Modifying the System Settings

Topics:

- System settings
- Modifying the PowerProtect Data Manager virtual machine disk settings
- Configure the DD system
- Virtual networks (VLANs)
- Syslog server disaster recovery
- Troubleshooting the syslog connection

System settings


You can use the PowerProtect Data Manager UI to modify system settings that are typically configured during PowerProtect Data Manager deployment.

To access **System Settings**, click .

Modify the network settings

Perform the following steps if you want to change the hostname or IP address of the PowerProtect Data Manager appliance, or modify other network settings such as the subnet mask, gateway, or DNS servers.

About this task

 **CAUTION:** Changing the hostname or IP address of the PowerProtect Data Manager appliance can require further actions to ensure the continued operation of external components. For more information, see [Changing the hostname or IP address](#).

Steps

1. From the PowerProtect Data Manager UI, click , and then click **Default Network**.
2. Update the following fields as necessary:
 - **Hostname**
 - **Primary DNS**
 - **Secondary DNS**
3. In the **Configuration Details** pane, click **Edit**, and then update the following fields for the IP address as necessary:
 - **IP Address**
 - **Subnet Mask**
 - **Gateway**
4. Click **Save**.

Changing the hostname or IP address

Changing the hostname or IP address of the PowerProtect Data Manager appliance can affect registered application hosts and VM Direct Engines.

If the PowerProtect Data Manager IP address has changed and you are using the File System agent, you must reregister the File System agent host with PowerProtect Data Manager using the new IP address. Follow the steps documented in the

PowerProtect Data Manager File System User Guide for reregistering the agent after a PowerProtect Data Manager IP address change.

If a VM Direct Engine is deployed for VMware virtual machine, Tanzu Kubernetes, or NAS protection, redeploy the protection engine. The *PowerProtect Data Manager Virtual Machine User Guide* provides instructions.

Modify the DNS search domain

Perform the following steps if you want to change the DNS search domain of the PowerProtect Data Manager appliance.

About this task

PowerProtect Data Manager automatically configures a search domain that is based on the domain name of the appliance. For example, if the FQDN of PowerProtect Data Manager is `ppdm.subdomain.domain.com`, the search domain is configured as `subdomain.domain.com`. This value can be modified, and more than one search domain can be used.

Steps

1. Use `ssh` to log in to PowerProtect Data Manager.
2. Run the following commands:

```
cd /usr/local/brs/puppet/scripts
./search_domains.sh
```

3. Follow the prompts to provide the new search domain information.

The following example adds the search domain `domain2.com` to the existing search domain `subdomain.domain.com`:

```
Setting search domains.
Current search domains: subdomain.domain.com
Change search domains to: subdomain.domain.com domain2.com
Applying search domains to [subdomain.domain.com domain2.com], input root password to
continue
[sudo] password for root:
New search domains: subdomain.domain.com domain2.com
```

Synchronizing the time between PowerProtect Data Manager and other systems

The PowerProtect Data Manager system time is synchronized with the ESXi host system.

The PowerProtect Data Manager system time must match those of the systems it interfaces with or compliance checks fail. It is recommended that all systems be configured to use an NTP server.


NOTE: Times displayed in the UI can use the time zone of each browser or use a configurable time zone that applies to all access regardless of the local time zone. The PowerProtect Data Manager system might be in a different time zone than that displayed by the UI. All log-file entries use the UTC time zone except those entries that are related to client browser connections, which use the server time zone.

Modify the user-interface time zone, system time zone, and NTP server

Use this procedure to modify the time zones and NTP server.

Steps

1. From the PowerProtect Data Manager UI, click  and then click **Time Zone**.

2. From the **User Interface Time Zone** list, select the applicable user-interface time zone. If a specific time zone is set instead of using that of the web browser, that time zone overrides the time zone used by the web browser when displaying information in the user interface.
3. From the **Server Time Zone** list, select the applicable time zone used by PowerProtect Data Manager. This time zone is used in component communication.
4. (Optional) To modify the NTP server:
 - a. Click .
 - b. In **NTP Servers**, provide the hostname or IP address of an NTP server.
5. Click **Save**.

Encryption in-flight

Using Transport Layer Security (TLS), you can encrypt backup or restore data that is in transit for centralized and self-service operations with DD Boost encryption. Encryption in-flight is available for agent host assets, Kubernetes cluster assets, Network-attached storage (NAS) assets, PowerStore block volume assets, and VMware virtual machine assets only.

By default, PowerProtect Data Manager supports an encryption strength of **HIGH** and uses DD Boost anonymous authentication mode. The DD Boost encryption software uses the **ADH-AES256-SHA** cipher suite. The *DD Boost for OpenStorage Administration Guide* provides more information about the cipher suite for high encryption.

Encryption in-flight is enabled for new installations. You can enable or disable encryption in-flight in the PowerProtect Data Manager UI. Enabling encryption in-flight is strongly recommended for all installations.

The following table lists the workloads and operations that support encryption in-flight:



 **NOTE:** Refer to the agent user guides for more information about the supported centralized and self-service operations.

Table 38. Supported workloads

Workload	Centralized backup	Centralized restore	Self-service backup	Self-service restore
File System with Application Direct	Yes	Yes (image-level restore only)	Yes	Yes (image-level restore only)
Kubernetes cluster	Yes	Yes	N/A	Yes (from the most recent backup)
Microsoft SQL Server with Application Direct	Yes	Yes (database-level restore only)	Yes	Yes (database-level restore only)
Microsoft Exchange Server with Application Direct	Yes	N/A	Yes	Yes
NAS	Yes	Yes	N/A	N/A
Oracle with Application Direct	Yes	N/A	Yes	Yes
SAP HANA with Application Direct	Yes	N/A	Yes	Yes
Virtual machines	Yes	Yes	N/A	N/A
PowerStore	Yes	Yes	N/A	N/A

Enabling encryption in-flight imposes additional overhead. Backup and restore performance for any client could be affected by 5-20%.

PowerProtect Data Manager supports encryption in-flight for all supported DD Boost and DDOS versions. The most up-to-date software compatibility information for PowerProtect Data Manager is provided by the E-Lab Navigator.

 **NOTE:** You do not need to enable in-flight encryption on connected DD systems. If DD encryption settings exist, the higher setting takes precedence.

Enable backup and restore encryption

You can ensure that the backup and restore content is encrypted when read on the source, transmitted in encrypted form, and then decrypted before it is saved on the destination.


Prerequisites

Review the information in Encryption in-flight to learn more about encryption in-flight.

The encryption settings determine if data transfers are encrypted during backup and restore operations.


- For File System, Microsoft Exchange Server, Oracle, SAP HANA, and Network Attached Storage (NAS) workloads, backup and restore encryption is only supported for Application Direct hosts. For Microsoft SQL Server, backup and restore encryption is supported for Application Direct and VM Direct hosts.
- When you add a new host to PowerProtect Data Manager, host configuration pushes the backup and restore encryption settings to the host.
- Only hosts that have the same version of PowerProtect Data Manager application agents installed support the host configuration.

Steps

1. From the PowerProtect Data Manager UI, click , and then select **Security**.
The **Security** dialog box appears.
2. Click the **Backup/Restore Encryption** switch so it is enabled, and then click **Save**.

Next steps

The **Jobs > System Job** window of the PowerProtect Data Manager UI creates a job to enable protection encryption. This job pushes encryption in-flight settings to the hosts to be used for self-service operations. Within the system job, a host configuration job is created for each host. If an error occurs, you can retry the system job or individual host configuration job.

 **NOTE:** For centralized backup and restore operations, PowerProtect Data Manager sends the encryption in-flight settings to the application agents on the Application Direct hosts and network-attached storage (NAS).

You can disable backup and restore encryption by clicking the **Backup/Restore Encryption** switch. PowerProtect Data Manager creates a system job in the **Jobs > System Job** window to disable backup and restore encryption.


Enable replication encryption

You can ensure that replicated content is encrypted while in-flight to the destination storage, and then decrypted before it is saved on the destination storage.

About this task


The encryption settings on both the source and destination systems must match for successful replication. For example, if you enable replication encryption in PowerProtect Data Manager, enable the setting on both source and destination before you define replication objectives. If you enable replication encryption after you initially define replication objectives, any replication jobs that were initiated during the period when the source and destination encryption settings did not match will fail.

Steps

1. From the PowerProtect Data Manager UI, click , and then select **Security**.
The **Security** dialog box appears.
2. Click the **Replication Encryption** switch so it is enabled, and then click **Save**.

Next steps

The **Infrastructure > Storage** window of the PowerProtect Data Manager UI displays the replication encryption setting for all protection storage systems.

 **NOTE:** For protection storage systems with DDOS version 6.2 and earlier installed, the status might display as *Unknown*. DDOS version 6.3 and later supports authentication mode. DDOS versions earlier than version 6.3 support only anonymous

authentication mode. PowerProtect Data Manager supports only anonymous and two-way authentication modes. Ensure that both source and destination use the same authentication mode.

You can take additional steps on your PowerProtect Data Manager server to enable in-flight encryption on connected DD systems by using **DD System Manager**, as described in the *DDOS Administration Guide*.

Additional considerations

Review the following additional considerations for encryption in-flight.

To validate that encryption is in use, you can check the status of existing connections on the DD system by running the `ddboost show connections` command in the DD Boost CLI:

If a connection was established with encryption, the value in the **Encrypted** column is *Yes*.

If a client establishes a connection with encryption, and establishes another connection without encryption, the value in the **Encrypted** column is *Mixed*. This circumstance might occur for one of the following reasons:

- Encryption settings that are defined on a per-client basis remain in place for a while after the client has disconnected. If the client previously established a connection without encryption and then later established a connection with encryption, the value shows as *Mixed*.
- Encryption settings are not specified for the DD Boost connections that are created on the application agent. The individual agent user guides provide more information.

If encryption settings exist on the DD and are also enabled in PowerProtect Data Manager, the higher setting takes precedence. As a result, the **Encrypted** column always shows *Mixed* or *Yes*.

Server monitoring with syslog

The syslog system logging feature collects system log messages and writes them to a designated log file. You can configure the PowerProtect Data Manager server to send event information in syslog format.

PowerProtect Data Manager serves as a syslog client to send diagnostic and monitoring data to the syslog server. You can access this data to perform audits, monitoring, and troubleshooting tasks.

The syslog server firewall is configured to receive data from PowerProtect Data Manager using the required ports listed in the *PowerProtect Data Manager Security Configuration Guide*. If your syslog server uses a port that is not listed, open the corresponding port on the PowerProtect Data Manager system.

Refer to the *PowerProtect Data Manager Security Configuration Guide* for the following information:

- Port usage
- Instructions for modifying firewall rules to add custom ports

It is recommended that you configure the PowerProtect Data Manager system to use an NTP server. NTP configuration is required to synchronize the PowerProtect Data Manager system time with the syslog server.

The selected severity level applies to all selected components. You cannot apply independent severity levels to each component. For example, selecting *Critical* forwards critical messages from all selected components. An exception is when you select *OS Kernel* or *PPDM Alert and Audit*, the corresponding audit log is forwarded by default, regardless of the selected severity level.

If no log messages are transmitted during a 24-hour period, PowerProtect Data Manager generates an alert to check the PowerProtect Data Manager and syslog server connection to verify that there are no problems preventing the exchange of messages.

Configure the syslog server

Use the following procedure to enable the syslog server, change the syslog server, change which events are forwarded, and disable syslog forwarding.


Prerequisites

To use TLS for the syslog connection:

- Import the syslog server security certificate into PowerProtect Data Manager. The *PowerProtect Data Manager Security Configuration Guide* provides instructions.

- By default, PowerProtect Data Manager uses `anon` authentication. If your syslog server uses another form of authentication, contact Customer Support.

Steps

1. From the PowerProtect Data Manager UI, click , select **Logs**, and then click **Syslog**. The **Logs** window opens to the **Syslog** page.

To enable syslog forwarding:

2. Move the **Syslog Forwarding** slider to the right to enable syslog forwarding.
3. Provide the following information:
 - **IP Address / FQDN**—IP address or fully qualified domain name of the syslog server.
 - **Port**—Port number for PowerProtect Data Manager and syslog server communications.
 - **Protocol**—Protocol to use for communications (TLS, UDP, or TCP).
 - **Components**—Syslog message components.
 - **Severity Level**—Specify the scope of the messages to forward to the syslog server.


To change the syslog server:

4. From the PowerProtect Data Manager UI, click , select **Logs**, and then click **Syslog**. The **Logs** window opens to the **Syslog** page.
5. Change the following syslog configuration details:
 - **IP Address / FQDN**—IP address or fully qualified domain name of the syslog server.
 - **Port**—Port number for PowerProtect Data Manager and syslog server communications.
 - **Protocol**—Protocol to use for communications (TLS, UDP, or TCP).

To change which events are forwarded:

6. From the PowerProtect Data Manager UI, click , select **Logs**, and then click **Syslog**. The **Logs** window opens to the **Syslog** page.
7. Change the **Components** and **Severity Level**.

To disable syslog forwarding:

8. From the PowerProtect Data Manager UI, click , select **Logs**, and then click **Syslog**. The **Logs** window opens to the **Syslog** page.
9. Move the **Syslog Forwarding** slider to the left to disable syslog forwarding.

To apply the changes:

10. Click **Save**.

Next steps

Once the syslog configuration is complete, check the connection status. Go to **System Settings > Logs > Syslog** and verify that the syslog server connection status indicates **Connected**. If the syslog server is not connected, the status indicates **Not Connected**.

Additional system settings

Some system settings directly relate to the deployment and maintenance of PowerProtect Data Manager.

For detailed information about the following topics, see *System Maintenance*.

- Licensing PowerProtect Data Manager
- Specifying a PowerProtect Data Manager host

Modifying the PowerProtect Data Manager virtual machine disk settings

Follow the steps in this section, under the guidance and recommendations of Customer Support, to expand the size of the data disk and system disk.

Modify the data disk size

Follow these steps to expand the size of a data disk that is single partitioned and has the log partition is on the system disk.

Steps

1. Perform the following steps from the **vSphere Web Client**:
 - a. Right-click the VM Direct appliance and select **Shut Down Guest OS**.
 - b. After the power off completes, right-click the appliance and select **Edit Settings**. The **Edit Settings** window appears with the **Virtual Hardware** button selected.
 - c. Increase the provisioned size of Hard disk 2 to the desired size, and then click **OK**.
NOTE: You cannot decrease the provisioned size of the disk.
 - d. Right-click the VM Direct appliance and select **Power On**.
2. Perform the following steps from the appliance console, as the root user.
NOTE: If you use `ssh` to connect to the appliance, log in with the admin account, and then use the `su` command to change to the root account.
 - a. Reboot the appliance by typing `reboot`.
 - b. On the **GNU GRUB** menu, press `Esc` to edit the GNU GRUB menu.
 - c. In the edit screen, search for the line that starts with `linux`, and then add word `single` before the entry `splash=0`. The following figure provides an example of the edit screen with the updated text.

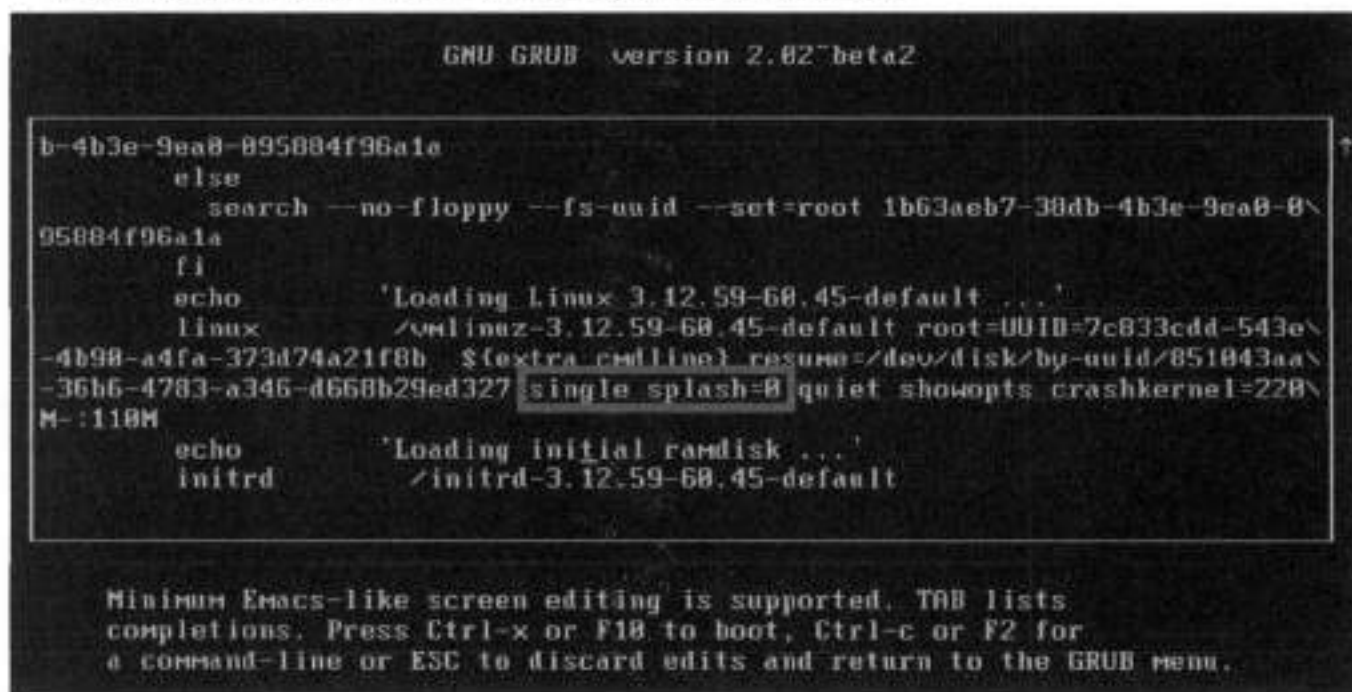


Figure 8. Editing the GNU GRUB menu

- d. Press **Ctrl-x** to reboot into single-user mode.

- e. When prompted, type the password for the root account.
- f. Unmount the data disk, by typing `umount /data01`.
- g. Start the partition utility, by typing `parted`, and then perform the following tasks:
 - i. Type `select /dev/sdb`.
 - ii. Type `print`. If you are prompted to fix issues, type `fix` at each prompt. The output displays the new disk size in the **Size** field and the current size in the table.
 - iii. Type `resize 1 new_size`. Where `new_size` is the value that appears in the **Size** field in the output of the `print` command.
 For example, to resize the disk to 700 GB, type: `resize 1 752GB`
 - iv. Type `quit`.

- 3. Reboot the VM Direct appliance by typing `systemctl reboot`.
- 4. Log in to the console as the root user.

NOTE: If you use `ssh` protocol to connect to the VM Direct appliance, log in with the admin account, and then use the `su` command to change to the root account.

- 5. Grow the xfs file system by typing `xfs_growfs -d /data01`.
- 6. Confirm the new partition size by typing `df -h`.

Modify the system disk size

Follow these steps to expand the size of a data disk when the log partition is the last partition on the system disk.

Steps

- 1. Perform the following steps from the **vSphere Web Client**:
 - a. Right-click the VM Direct appliance and select **Shut Down Guest OS**.
 - b. After the power off completes, right-click the appliance and select **Edit Settings**. The **Edit Settings** window appears with the **Virtual Hardware** button selected.
 - c. Increase the provisioned size of Hard disk 1 to the desired size, and then click **OK**.

NOTE: You cannot decrease the provisioned size of the disk.

 - d. Right-click the VM Direct appliance and select **Power On**.
- 2. Boot from a SuSE Linux Enterprise Server (SLES) version 12 CD.
- 3. Start the partition utility, by typing `parted`, and then perform the following tasks.
 - a. Type `select /dev/sdx`.
 - b. Type `print`. If you are prompted to fix issues, type `fix` at each prompt. The output displays the new disk size in the **Size** field and the current size in the table.
 - c. Type `quit`.
- 4. Reboot the VM Direct appliance by typing `systemctl reboot`.
- 5. Log in to the console as the root user.

NOTE: If you use `ssh` protocol to connect to the VM Direct appliance, log in with the admin account, and then use the `su` command to change to the root account.
- 6. Grow the xfs file system by typing `xfs_growfs -d /data01`.
- 7. Confirm the new partition size by typing `df -h`.

Configure the DD system

Prerequisites

Before you can use DD to protect the system, use NFS to export the MTree that PowerProtect Data Manager uses on the DD system. The setup on the DD system requires that you add the PowerProtect Data Manager client with `no_root_squash`.

Steps

1. Use a web browser to log in to the **DD System Manager** as the system administrator.
2. In the **Summary** tab, **Protocols** pane, select **NFS export > create export**.
The **Create NFS Exports** window appears.
3. In the **Create NFS Exports** window:
 - a. In the **Export Name** field, specify the name of the DD MTree.
 - b. If you have not yet created the DD MTree, follow the prompts to create the MTree and click **Close**.
 - c. In the **Directory path** field, specify the full directory path for DD MTree that you created. Ensure that you use the same name for the directory.
 - d. Click **OK**.
A message appears to indicate that the NFS export configuration save is in progress and then complete.
 - e. Click **Close**.

Virtual networks (VLANs)

PowerProtect Data Manager can separate management and backup traffic onto different virtual networks (VLANs). Virtual networks help to improve data traffic routing, security, and organization.

The default configuration routes the management traffic over the same network as backup traffic. All assets are part of the same network.




Figure 9. Flat network

You can also configure virtual networks to separate management traffic from backup traffic. This configuration can also separate traffic that originates from different networks. In that case, you can use the same virtual network for management and backup traffic, or separate virtual networks for each.




Figure 10. Virtual networks

To use virtual networks with PowerProtect Data Manager, you must configure the DD and network infrastructure before you configure the PowerProtect Data Manager or assign networks to assets.

Configuration follows a multistep workflow:

1. Configure the virtual network on the DD.
2. Add the DD as storage and name the network interface.
3. Add the virtual network to the PowerProtect Data Manager.
4. Register the assets with the PowerProtect Data Manager.
5. Create a protection policy (or edit an existing policy) and assign the preferred virtual network.
6. Optionally, assign the virtual network to individual assets. This action overrides any preferred virtual network that you may have specified through a protection policy.

The initial steps to configure and add each virtual network are one-time events. The subsequent steps to assign virtual networks to protection policies or assets happen as required.

Configuration is nondisruptive. You can add, edit, or delete virtual networks without affecting background activities, disconnecting network interfaces, or affecting the PowerProtect Data Manager user interface.

PowerProtect Data Manager logs network changes in the audit log. Failed network changes appear in the **System** alerts.

Virtual network traffic types

PowerProtect Data Manager supports virtual networks for the following traffic types:

Table 39. Traffic types

Type	Description
Management	Control traffic, typically HTTPS REST API operations; small file transfers, such as logs and update packages; other essential traffic, such as identity provider authentication.
Data	Large amounts of customer data, such as backup and restore traffic, cloud tiering, and CloudDR traffic.
Data for Management Components	Customer data that is related to management and control operations, such as ServerDR, indexing and searching, replication monitoring, and copy deletion.

The Data for Management Components type carries traffic which relates to management operations but which can contain customer information. Where required, you can separate this traffic from either the Management network, the Data network, or both.

For example, some environments may support different speeds for each network: a 1 Gbps network for management and a 10 Gbps network for data. Other environments may have policies or rules that govern whether customer data can flow across the Management network. Separating the Data for Management Components traffic enables you to optimize flow for security, speed, and other priorities.

Virtual network planning

When you plan your virtual network configuration, observe the following requirements:

Table 40. Component traffic type requirements

Component	Compatible types	Incompatible types
PowerProtect Data Manager	Management, Data for Management Components	Data
Protection engines	Data for Management Components, Data	Management
Search Engine nodes	Data for Management Components	Management, Data
Reporting Engine	Management, Data for Management Components	Data

While the table indicates compatible traffic types, protection engines can operate without virtual networks.

Separating the Data for Management Components traffic from the Management traffic requires you to name the virtual networks for protection storage. Change network settings for protection storage provides instructions. If you do not name the virtual networks for protection storage, this traffic defaults to the Management network.

Parallel virtual networks

Your environment may have more than one virtual network for each traffic type, such as different Data networks for different departments. Where parallel virtual networks exist, all protection engines require an interface to at least one virtual network of each required type. However, each protection engine does not require connections to all virtual networks of the required types.

For example:

- Your environment has Finance and Engineering departments with their own assets.
- Your environment has the following virtual networks: Management, Finance Data, and Engineering Data.

The following table describes the connections to each virtual network for scenarios where both departments share a protection engine and where departments have private protection engines.

Table 41. Example: virtual network interfaces

Virtual network name	Shared protection engine	Private protection engines	
		Finance protection engine	Engineering protection engine
Management	Yes	Yes	Yes
Finance Data	Yes	Yes	No
Engineering Data	Yes	No	Yes

Even though protection engines require connections for Data traffic, the private protection engines maintain separation between the virtual networks for each department.

Several of the diagrams for supported virtual network topologies include parallel virtual networks.

Virtual network topologies

The following diagrams illustrate the supported virtual network topologies and how they relate to traffic types:

Single network

This topology assigns all traffic types to the same network. There is no separation between Management and Data or between agents which belong to different logical organizations.

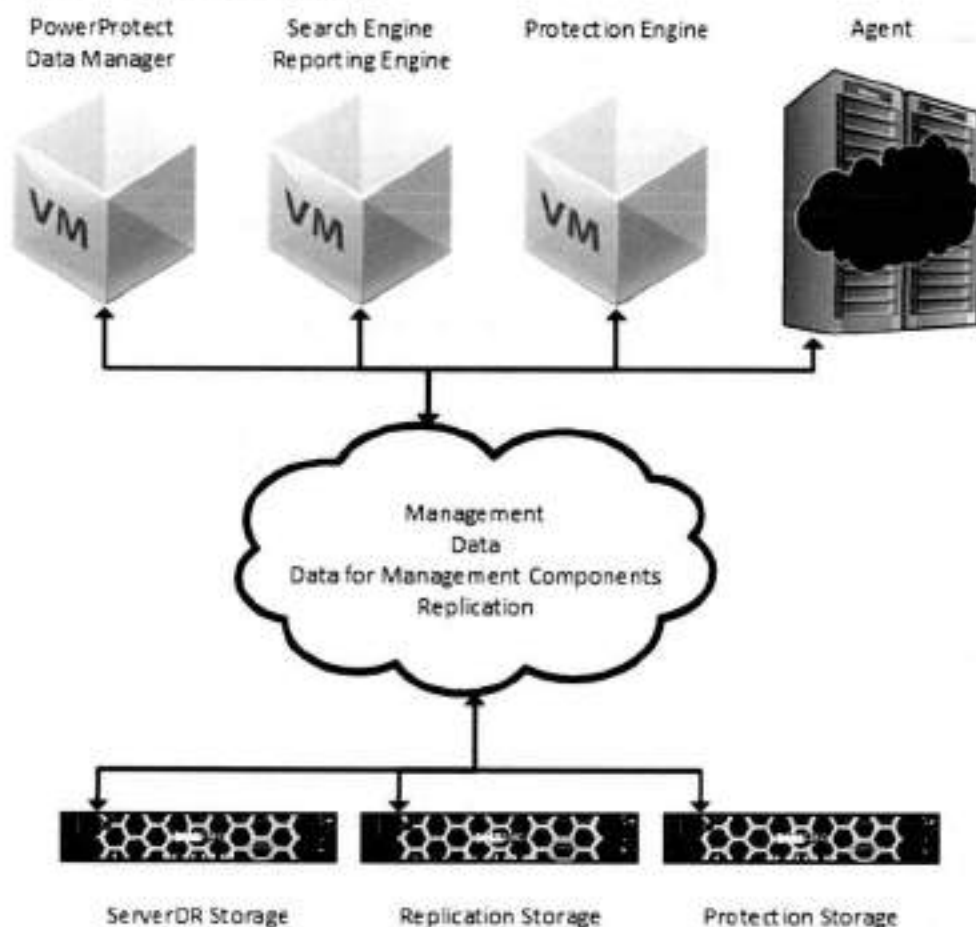


Figure 11. Single network

Data for Management Components traffic on Management network

This topology separates Management traffic from Data traffic but keeps the Data for Management Components traffic with the Management traffic.

This tradeoff operates well in environments where the Management network can support frequent large data transfers and which allow customer data on the Management network.

Thick lines indicate paths that transfer comparatively more data, such as files and update packages. Thin lines indicate paths that transfer comparatively less data, such as HTTPS API traffic only.

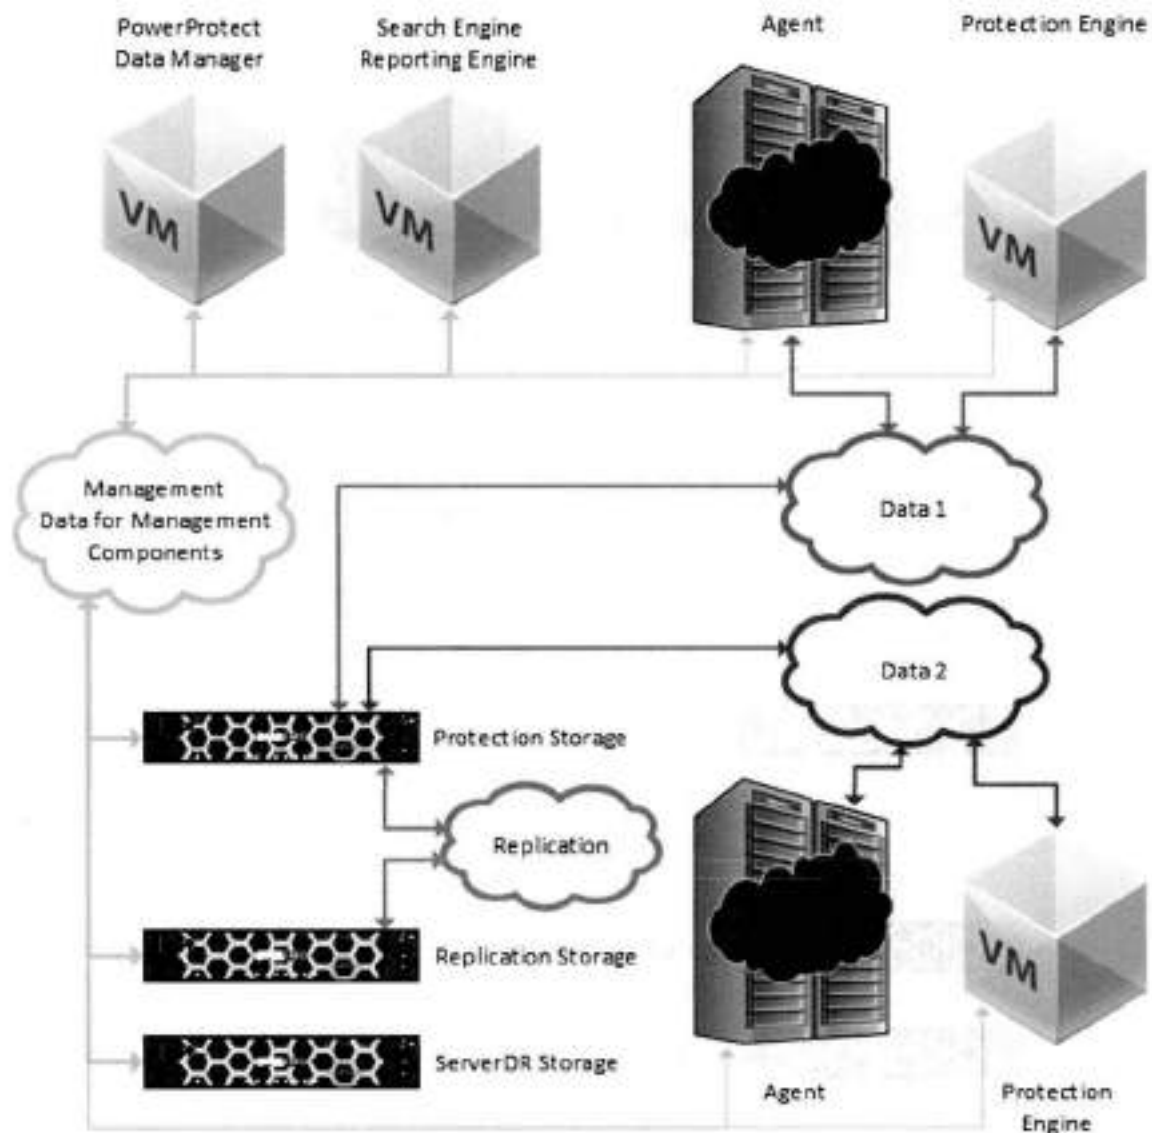


Figure 12. Data for Management Components traffic on Management network

Data for Management Components traffic on Data network

This topology separates Management traffic from Data traffic but keeps the Data for Management Components traffic with the Data traffic.

This tradeoff operates well in environments where the Management network cannot support frequent large transfers or which do not allow customer data on the Management network. However, there is no separation between backup data and control data, and Data for Management Components traffic competes with other traffic.

Thick lines indicate paths that transfer comparatively more data, such as files and update packages. Thin lines indicate paths that transfer comparatively less data, such as HTTPS API traffic only.

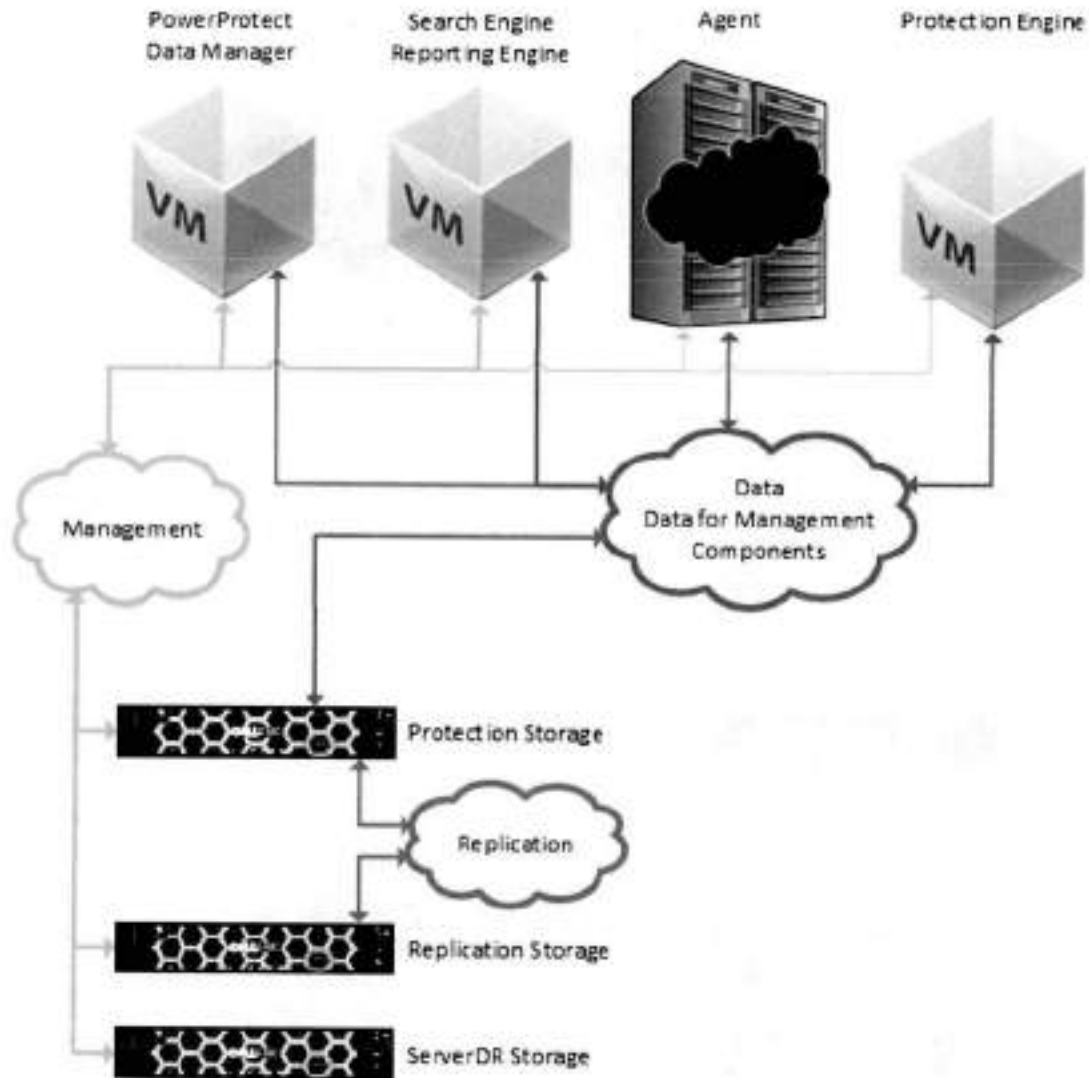


Figure 13. Data for Management Components traffic on Data network

Full separation

This topology implements complete separation between all traffic types for maximum throughput and security. Customer data does not flow across the Management network.

Thick lines indicate paths that transfer comparatively more data, such as files and update packages. Thin lines indicate paths that transfer comparatively less data, such as HTTPS API traffic only.

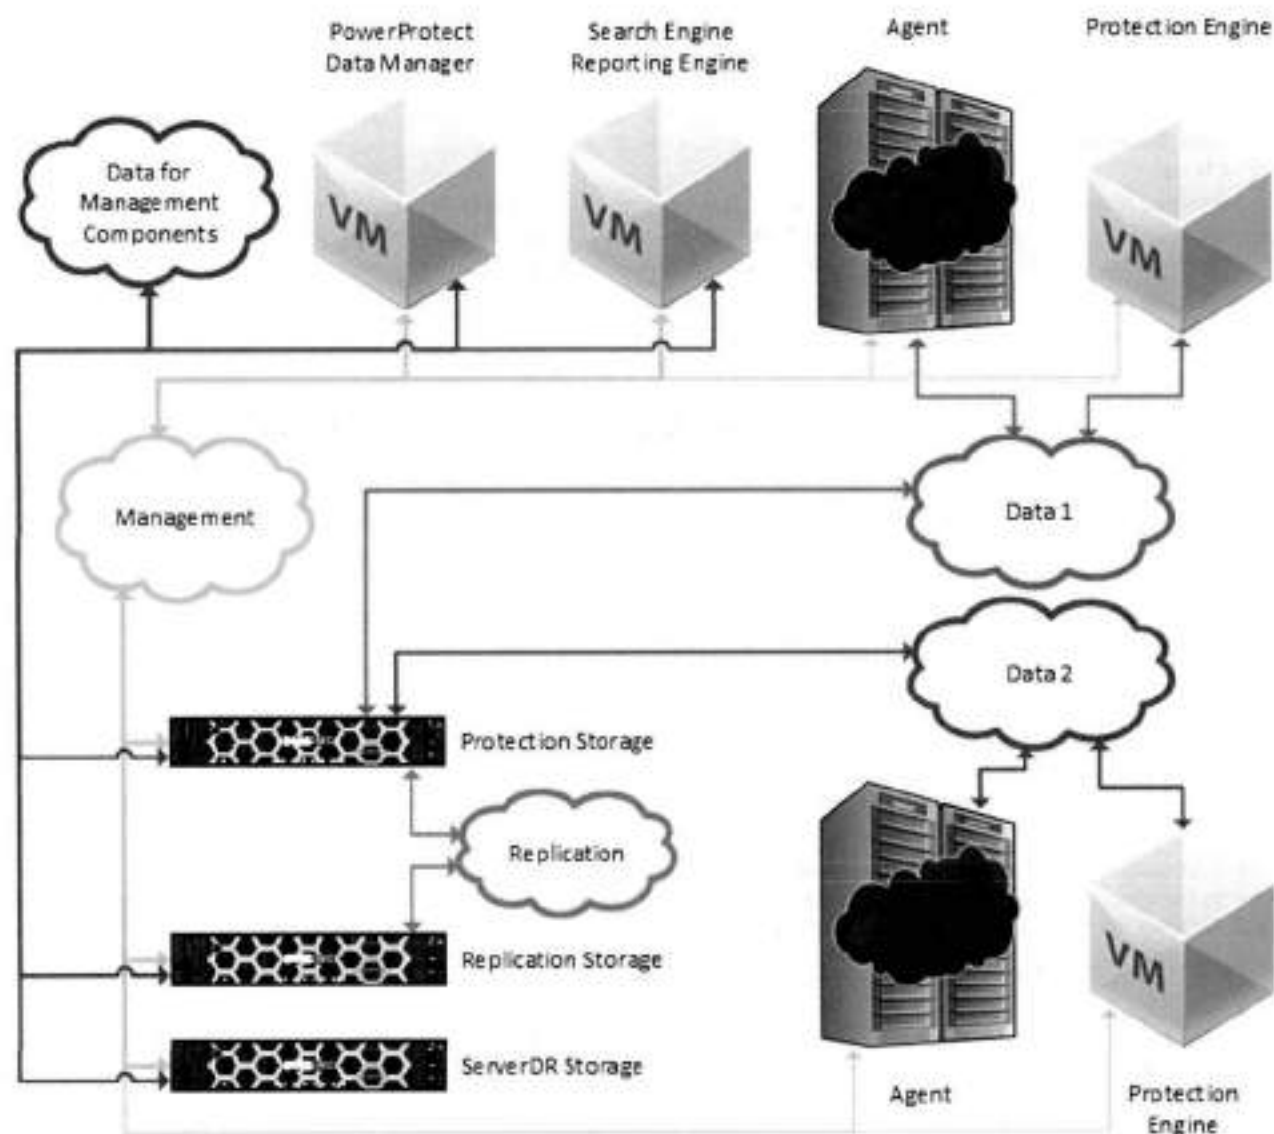


Figure 14. Full separation

Supported scenarios

PowerProtect Data Manager supports virtual networks for the following use cases:

- Virtual machine backups
- Kubernetes backups
- Database backups
- Microsoft Exchange Server backups
- File system backups
- Replication
- Disaster recovery
- Cloud DR
- Storage Data Management
- Search Engine

NOTE: The first time that you use the **Networks** page to add a virtual network to an environment with existing Search Engine nodes, PowerProtect Data Manager does not automatically add the virtual network to the Search Engine. Instead, manually edit each Search Engine node to add the virtual network. This action makes the Search Engine aware of virtual networks. Any subsequent new virtual networks are automatically added to the Search Engine.

Virtual network prerequisites

Before you configure a virtual network, complete the following actions:

- Register the vCenter server on which PowerProtect Data Manager is deployed. You can verify this on the **vCenter** tab of the **Asset Sources** page. You can also add a hosting vCenter. Specifying the PowerProtect Data Manager host provides instructions.
- Configure the network switch port for trunk mode. This setting allows the port to carry traffic for multiple VLANs.
- Enable Virtual Guest Tagging (VGT) or Virtual Switch Tagging (VST) mode on the VMware ESXi virtual network switch port for PowerProtect Data Manager. You can use a standard port group or a distributed port group.
 - VGT—For port groups on standard virtual switches, configure the virtual switch port for VLAN ID 4095, which makes all VLANs accessible. For port groups on distributed virtual switches, use VLAN trunking, which supports specifying multiple VLANs by ID or range. For more information, see the VMware ESXi documentation.
 - VST—You can configure the port group with a VLAN ID from 1-4094.
- Configure a VLAN interface for the DD through the **Interfaces** tab on the **Hardware > Ethernet** window in the DD System Manager. The DD documentation provides more information.

It is recommended that you choose an interface name that incorporates the VLAN ID. For example, the interface name `ethV1.850` for VLAN ID 850.

- Add the DD as protection storage for PowerProtect Data Manager.

PowerProtect Data Manager does not verify the network switch configurations. If the physical or virtual network switch is incorrectly configured, then virtual network configuration fails.

Configuring virtual networks


The following topics create and maintain virtual networks in PowerProtect Data Manager for use with assets on different VLANs.

PowerProtect Data Manager names each virtual network in two places: the interface to the protection storage system and the interface to the protected assets. These names are not required to match. However, it is strongly recommended that you use the same network name in both locations for each virtual network. Record each network name for later use.

It is also recommended that you choose network names that incorporate the VLAN ID. For example, `sales-vlan850` for VLAN ID 850.

Adding a virtual network includes creating a pool of static IP addresses. PowerProtect Data Manager uses these addresses for the local interfaces to the virtual network and for any VM Direct protection engines or Search Engine nodes that you deploy on this network.

Each VM Direct protection engine or Search Engine node requires an IP address on the virtual network. The PowerProtect Data Manager interface requires one IP address. Ensure that you have enough IP addresses available on each network to meet this requirement. To prepare for future expansion, you can add more IP addresses than are initially required.

When you review the list of virtual networks, rows that require attention are indicated with a  beside the name. View the network details for more information.

Add a virtual network

Configure a new virtual network for use with assets and protection policies.



About this task

Each new virtual network requires at least one IP address for each PowerProtect Data Manager network interface. Review the **Number of IP addresses needed** field before you supply the required static IP addresses.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Networks**. The **Networks** window appears.
2. Click **Add**. The **Add Network** wizard opens.
3. For **Purpose**, select one or more traffic types.

Virtual network traffic types provides more information.

4. In the **Network Name** field, type the name of the new virtual network.
It is recommended that you keep the network names consistent for each VLAN.
5. In the **VLAN ID** field, type the numeric value 1 through 4094 that corresponds to the VLAN which this virtual network represents.
6. Provide the **MTU** (maximum transfer unit) for the virtual network.
Allowable **MTU** values range from 1500 to 9000.
7. Click **Next**.
The **Add Network** wizard moves to the **Static IP Pool** page.
8. From the **Static IP Pool** page:
 - a. Select the **Type** of IP pool.
If you need more than one type of IP pool, click **Add Alternate Configuration Details**. You can edit this additional IP pool by clicking **Edit**, or delete it by clicking **Delete**.
 - b. Provide the **Subnet Mask** for an IPv4 pool or the **Prefix** for an IPv6 pool.
 - c. Provide the number of reserved IP addresses for PowerProtect Data Manager to use for communication on this virtual network.
You can add or remove individual IP addresses or ranges of IP addresses.
 - To add an individual IP address or range of IP addresses, click , select **Value** or **Range**, and then provide the value or range.
 - To remove an individual IP address or range of IP addresses, click  next to its entry.
9. Verify that the static IP address pool contains enough addresses to add the virtual network.
10. Click **Next**.
The **Add Network** wizard moves to the **Routes** page.
11. If applicable, click **Add** to define any required routes.
The **Add Routes** page opens. Complete the following substeps:
 - a. Select a route type:
 - If you select **Subnet**, define the subnet in CIDR format. For example, 10.0.0.0/24 for IPv4 or fe80:7f03:79a5:2d11::f9a5/64 for IPv6.
 - If you select **Host**, type the IP address.
 - b. Type the IP address of the default gateway through which PowerProtect Data Manager should reach the subnet or host.
 - c. Click **Add**.
The **Add Routes** page closes. The **Routes** list displays the new route.
 - d. Review the route information.
If any parameters are incorrect, select the checkbox for that route and then click **Delete**.
 - e. Repeat these substeps for any additional required routes.
12. Click **Next**.
The **Add Network** wizard moves to the **Summary** page.
13. Verify the network configuration information, and then click **Finish**.
The **Add Network** wizard closes. The **Networks** page displays the new network with the **Initiating** status.

Next steps

PowerProtect Data Manager may take a short time to configure the virtual network.

If the virtual network status changes to **Failed**, then a corresponding system alert contains more information about the cause of the failure. Troubleshoot the failure and then complete one of the following actions:

- If the failure was caused by a configuration issue, click **Edit** to update the network configuration.
- If the failure was transient or had an external cause, and the configuration is correct, click **Retry** to use the same settings.

NOTE:


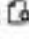
When you edit or retry a virtual network operation that failed and there are additional IP addresses in the address pool, PowerProtect Data Manager marks the last failed IP address as abandoned. PowerProtect Data Manager does not try to reuse any IP addresses that are marked as abandoned. The UI does not display this condition.

KB article 000181120 provides more information about how to use the REST API to detect when an IP address is marked as abandoned. The article also provides steps to correct this condition so that the IP address can be used again.

View the details of a virtual network

If the virtual network name is ambiguous, you can view the details to further identify the virtual network before making changes. You can also identify components that require attention after a change.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Networks**.
The **Networks** window appears.
2. Locate the row that corresponds to the appropriate virtual network.
The columns for each row indicate the associated VLAN ID and network status. Rows that require attention are indicated with a  beside the name.
3. Click  for that row.
The **Details** pane opens to the right.
This pane contains information about the virtual network configuration, such as the static IP address pool details, assigned traffic types, and configured routes. This pane also lists any components that are configured with an interface on this network, their types, and their assigned IP addresses.
4. Click **X** to close the details pane.

Changing virtual network traffic types after configuration

Under normal operation, you configure a virtual network and then assign the interface to new or existing components which support the selected traffic types. However, should your environment change, you can later change the traffic type settings for a virtual network.

After you reconfigure a virtual network, the new traffic type settings may no longer align with the interface assignments for existing components on that virtual network. In these cases, PowerProtect Data Manager notifies you about conflicts between traffic types and interface assignments, but does not take automatic action.

Instead, the UI marks conflicts with a warning symbol (). Administrators should review any warnings and edit the indicated components to manually remove the incompatible network interfaces. For example:

- Search Engine node interfaces to virtual networks that carry Data traffic, but not Data for Management Components traffic.
- Protection engine interfaces to virtual networks that carry Data for Management Components traffic.
- PowerProtect Data Manager interfaces to virtual networks that carry Data traffic, but not Data for Management Components traffic.

Under these circumstances, PowerProtect Data Manager continues to operate normally. However, resolving the conflict returns the IP address to the address pool.

Edit a virtual network

You can change any parameter for a virtual network without deleting the network. For example, to add more IP addresses to the static IP pool.

Prerequisites

If an IP address from the static IP pool is already in use, you cannot remove the address from the pool.

Before you change the traffic types for a network, disable indexing. Set up and manage indexing provides instructions.

About this task

After deployment, the default network has all traffic types enabled. You can remove the Data and Data for Management Components types from this network, but not the Management type.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Networks**.
The **Networks** window appears.
2. Locate the row that corresponds to the appropriate virtual network, and then click the radio button to select that row.
The PowerProtect Data Manager enables the **Edit** and **Delete** buttons.
3. Click **Edit**.
The **Edit Network** wizard opens to the **Summary** page.
4. Click **Edit** for the **Configuration**, **Static IP Pool**, and **Routes** sections.
The **Edit Network** wizard moves to the **Configuration**, **Static IP Pool**, or **Routes** page.
5. Modify the appropriate network parameters, and then click **Next**.
If you modify the virtual network in a way that requires more IP addresses, you cannot continue until you add more addresses to the static IP address pool.
The **Edit Network** wizard moves to the **Summary** page.
6. Verify the network configuration information, and then click **Finish**.
The **Edit Network** wizard closes. The **Networks** page reflects the updated information, where applicable.
You may need to view the details for the virtual network to verify some changes.

Next steps

If you disabled indexing, re-enable indexing. Set up and manage indexing provides instructions.

Delete a virtual network

Although optional, it is recommended that you delete virtual networks when they are no longer required.

Prerequisites

- Unassign the virtual network from any applicable assets.
- Disable indexing. Set up and manage indexing provides instructions.
- Disable every VM Direct Engine that is configured to use the virtual network.
- Disable every Search cluster that uses the virtual network.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Networks**.
The **Networks** window appears.
2. Locate the row that corresponds to the appropriate virtual network, and then click the radio button to select that row.
PowerProtect Data Manager enables the **Edit** and **Delete** buttons.
3. Click **Delete**.
4. Verify the network information, and then click **OK** to acknowledge the deletion warning.
The PowerProtect Data Manager removes the virtual network from the list on the **Networks** page.

Next steps

Re-enable indexing, VM Direct Engines, and Search clusters.

Change network settings for protection storage

After you add protection storage, name the virtual network or networks between the PowerProtect Data Manager and the protection storage system. To rename a virtual network (edit the network name), repeat these steps.

About this task

Separating the Data for Management Components traffic from the Management traffic requires you to name the virtual networks for protection storage. If you do not name the virtual networks for protection storage, components such as PowerProtect Data Manager and Search Engine nodes have no route to protection storage over the Data for Management Components network. This traffic defaults to the Management network.

NOTE: Network interfaces that exist on a DD 7.4.x or earlier system and that are configured to use an uncompressed IPv6 format cannot be discovered. An example of an uncompressed IPv6 format is 2620:0000:0170:0597:0000:0000:0001:001a. An example of a compressed IPv6 format is 2620:0:170:597::1:1a. To use these network interfaces, reconfigure them to use either an IPv4 address or a compressed IPv6 address, and then initiate a discovery.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Storage**. The **Storage** window appears.
2. On the **Protection Storage** tab, select the storage system, and then select **More Actions > Change Network Settings**. The **Change Network Settings** window opens and displays a list of known network interfaces, assigned IP addresses, link speeds, and network purposes.
3. Identify the interfaces for each new virtual network, and then select or type names for the virtual networks in the corresponding fields.
Each interface indicates an IP address, link speed, and network purposes.
4. If you typed a name for a virtual network in step 3, select one or more network purposes for the virtual network.
5. Click **Save**.
The PowerProtect Data Manager stores the network names.

Virtual network asset assignment

Assignments identify which assets should use each virtual network. There are two methods to associate an asset with a virtual network:

- By protection policy

You can configure the PowerProtect Data Manager to choose a preferred virtual network for all assets on a protection policy.

- By asset

You can assign virtual networks to individual assets. This method is optional and overrides any virtual network assignment from a protection policy. Assets which are not individually assigned automatically use the preferred virtual network.

You can use this method to specify a virtual network for any asset. However, this method is especially suited to configuring assets which are exceptions to the rule. You can also split assets on the same application host across multiple virtual networks. For example, when an asset has its own network interface or belongs to another department.

It is recommended that you assign assets to virtual networks by protection policy, where possible.

Before you assign an asset, perform the following actions:

- Test connectivity from the asset host to the PowerProtect Data Manager by pinging the PowerProtect Data Manager IP address on that virtual network.
- Register the asset source with the PowerProtect Data Manager.
- Approve the asset source.

Assign a virtual network by protection policy

The following steps apply a virtual network to an existing protection policy. You can also assign a virtual network when you create a protection policy.

About this task

The **Network Interface** field selects the network interface for communication with the destination protection storage system. This network carries the backup data.

Steps

1. From the PowerProtect Data Manager UI, select **Protection > Protection Policies**. The **Protection Policies** window appears.
2. Locate an existing protection policy for which you want to configure a virtual network.

3. Select the radio button for the protection policy, and then click **Edit**.
The **Edit Policy** wizard opens to the **Summary** page.
4. In the **Objectives** block, click **Edit**.
The **Edit Policy** wizard moves to the **Objectives** page.
5. Select the checkbox for the appropriate schedule.
6. In the **Network Interface** field, select the correct virtual network from the list. Only network interfaces with a network purpose of **Data** are listed. Review the section *Change network settings for protection storage when changing network settings*.

Each list entry indicates the interface name, interface speed, and virtual network name.

If the network was not named, a combination of the interface name and VLAN ID replaces the virtual network name. For example, `ethV1.850`. An interface without a virtual network name behaves as if a virtual network was not configured.
7. Click **Next**.
The **Edit Policy** wizard moves to the **Summary** page.
8. Verify the policy information, and then click **Finish**.
Ensure that the selected assets are part of the virtual network.
The **Edit Policy** wizard closes.
9. Click **OK** to acknowledge the update, or click **Go to Jobs** to monitor the update.

Assign a virtual network by asset

This procedure is optional. You can assign a virtual network for individual assets or for all assets on a particular application host.

About this task

This setting overrides the network assignment from the protection policy. If PowerProtect Data Manager cannot use this network assignment for any reason, the setting falls back to the assignment from the protection policy.

- NOTE:** You cannot back up individual assets across different networks on the same protection policy and application host. Instead, create a separate protection policy for the assets on each network.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
The **Assets** window appears.
2. Locate the appropriate assets from the list on any tab.
Use the checkbox to select each asset. You can select more than one asset at a time.
3. Click **More Actions > Assign Network**.
The **Associated Assets** window opens.
4. To use the virtual network for all assets on the same application host, click **Include**.
Otherwise, to use the virtual network for only the selected assets, click **Do Not Include**. Consider whether you require a separate protection policy for assets on different networks.
The **Assign Network** window opens.
5. Select a virtual network from the **Network Label** list, and then click **Save**.

Results

The PowerProtect Data Manager applies the network selection to the selected assets. The **Network** column in the list of assets for each tab now indicates the selected virtual network.

Syslog server disaster recovery

Use the following procedure to restart the log manager service for the syslog server in a server disaster recovery (DR) scenario.

Prerequisites

After disaster recovery of the PowerProtect Data Manager system is complete, perform the following steps on the restored PowerProtect Data Manager system.

Steps

1. Verify that all PowerProtect Data Manager services are running in **System Settings > Support > System Services Status**.
2. Restart the `logmgr` service by running the command `logmgr restart`, and then wait for a few seconds for the service to restart.

Next steps

If your syslog server uses a custom port, open the corresponding port on the restored PowerProtect Data Manager system. The *PowerProtect Data Manager Security Configuration Guide* provides more information.

Troubleshooting the syslog connection

Review the following information that is related to troubleshooting the syslog connection.

No messages are transmitted to the syslog server

Log messages are generated in the PowerProtect Data Manager services log files, however these messages are not transmitted to the syslog server. If this issue occurs, complete the following tasks:

1. Verify that the PowerProtect Data Manager firewall is using the required ports. If your syslog server uses a different port, open the corresponding port on the PowerProtect Data Manager system.
2. Verify the syslog server firewall. Ensure that the ports are configured to accept data.
3. Verify that the protocol is the same for both PowerProtect Data Manager and the syslog server. If you are using TLS, PowerProtect Data Manager uses `anon` authentication by default. If your syslog server uses another form of authentication, contact Customer Support.

Managing Reports

Topics:

- PowerProtect Data Manager reporting
- Port requirements
- Server requirements
- Unsupported reporting engine vCenter operations
- Known issues with the reporting engine and Report Browser
- Configure and deploy the reporting engine
- Updating the reporting engine from version 19.10
- Report Browser
- Deleting the reporting engine
- Managing disaster recovery of the reporting engine

PowerProtect Data Manager reporting

PowerProtect Data Manager comes with a reporting engine that offers reporting capabilities from within the PowerProtect Data Manager user interface. You can access built-in report templates that you can directly run to generate reports. Feedback can be provided for future releases.

These reports help you retrieve information about the data protection activities in your environment. Using these reports, you can diagnose problems, plan to mitigate risks, and forecast future trends. You can also run reports on-demand and export reports in CSV format.

All events in report data are shown in UTC.

PowerProtect Data Manager reporting is available for on-premises PowerProtect Data Manager deployments.

NOTE: PowerProtect Data Manager reporting is not supported with PowerProtect Data Manager in cloud environments.

Configure the reporting engine to set up reporting capabilities for PowerProtect Data Manager. After the reporting engine is configured, you can run reports from **Reports > Report Browser**.

NOTE: If you are using another reporting tool such as CloudIQ, you can choose not to configure PowerProtect Data Manager reporting.

Port requirements

The following table summarizes the port requirements for PowerProtect Data Manager and the reporting engine. The *PowerProtect Data Manager Security Configuration Guide* provides more information about ports for PowerProtect Data Manager. Read this table in conjunction with the port usage topic for PowerProtect Data Manager.

Table 42. Reporting engine port requirements

Source system	Destination system	Port	Protocol	TLS supported	Notes
PowerProtect Data Manager	Reporting engine	9002	TCP	TLS 1.2	REST API service.
PowerProtect Data Manager	Reporting engine	9613d	Proprietary	TLS 1.2	Infrastructure node agent management of the reporting engine.

Table 42. Reporting engine port requirements (continued)

Source system	Destination system	Port	Protocol	TLS supported	Notes
Reporting engine	PowerProtect Data Manager	8443	TCP	TLS 1.2	REST API service for collecting reporting data.
User	Reporting engine	22	SSH	TLS 1.2	SSH for support and administration. Encrypted by private key or optional certificates.

Server requirements

Observe the following requirements for the reporting engine.

- SUSE Linux Enterprise Server (SLES) version 12 SP5
- 8 vCPUs, 16 GB RAM
- Disk 01: 48 GB to install the operating system and Reporting Application Server
- Disk 02: 512 GB to store the reporting data
- Disk 03: 8 GB to store log information

NOTE: The reporting engine only supports IPv4 communication.

Unsupported reporting engine vCenter operations

The reporting engine should only be managed according to documentation and guidance.

Unless communicated by Customer Support, using vCenter to change or control the virtual machine where the reporting engine is deployed is unsupported. Unsupported vCenter operations include:

- Changing the virtual machine power state
- Changing the virtual machine properties
- Cloning the virtual machine
- Deleting the virtual machine
- Performing a manual vMotion
- Taking or restoring snapshots

Known issues with the reporting engine and Report Browser



Administrators should familiarize themselves with the known issues of the new reporting feature before using it. Understanding the known issues will help with the maintenance of the feature and interpretation of the reports.

The following table describes the known issues of the new reporting feature.

Table 43. Known issues with the reporting engine and Report Browser

Issue
<p>You might receive error messages similar to the following when trying to configure the Report Browser or generate a report:</p> <pre>The reporting engine is not configured Configure the reporting engine to access reports. An error occurred in obtaining the reporting engine configuration details.</pre> <p>To resolve this issue:</p> <ol style="list-style-type: none"> 1. If the Report Browser has not been configured, configure it.

Table 43. Known issues with the reporting engine and Report Browser (continued)

Issue
<p>2. If the Report Browser has been configured or you are trying to configure it, reload the dashboard in your web browser.</p> <p> NOTE: PowerProtect Data Manager only supports the latest version of Google Chrome.</p>
<p>The Report Browser is not integrated with Cloud Snapshot Manager. Since the Report Browser does not display Cloud Snapshot Manager jobs, it might display a lower total job count than the total job count displayed in the Protection Jobs window.</p>
<p>The Report Browser displays an entry for each retried job that originally had a Failed status. This display of multiple entries does not match the behavior of the Protection Jobs window, which only displays a single entry for the failed job. This discrepancy can result in the Report Browser displaying a higher total job count than the total job count displayed in the Protection Jobs window.</p>
<p>If the total amount of data transferred for a job is less than 1 MB, the job entry shows 0 bytes in the Data Transferred column.</p>
<p>The reporting engine is listed as an unidentified entry in the Application Agents pane.</p> <p> CAUTION: Do not remove this entry. If you do, see Configure and deploy the reporting engine.</p>
<p>The selection of SMS assets from a custom filter is ignored. Even if these assets are selected, they are not displayed on the report that uses the filter.</p>
<p>In the Jobs Summary - Table View report, the search functionality only supports an "equals" filter type.</p>
<p>If you edit an existing custom scope, using the search function removes the previously selected assets. To add new assets and keep the previously selected assets, do not use the search function. Add new assets by scrolling through the list of all assets and changing the current selections.</p>

Configure and deploy the reporting engine

Perform the following steps in the PowerProtect Data Manager UI to configure and deploy the reporting engine.

Prerequisites

- You must deploy the reporting engine on a separate virtual machine.
- The vCenter server must be added as an asset source from **Infrastructure > Asset Sources**.
- The virtual machine requires 500 GB to function properly.

About this task

It is recommended that you deploy the reporting engine to the vCenter server that hosts PowerProtect Data Manager. To verify the hosting vCenter:

1. Click the **Settings > Hosting vCenter** link.
2. Provide the details for the vCenter server that hosts PowerProtect Data Manager or select the hosting vCenter server from asset sources.

Steps

1. From the PowerProtect Data Manager UI, select **Reports > Reporting Engine**.
2. Click **Configure**.
The **Configure Reporting Engine** dialog box opens.
3. In the **Configure Reporting Engine** dialog box, complete the required fields:
 - **vCenter server to deploy**—Specify the vCenter server on which to deploy the reporting engine.
If you specified the hosting vCenter server, PowerProtect Data Manager populates the fields with the required information.
 - **ESX host or cluster**—Select on which cluster or ESXi host you want to configure the reporting engine.
 - **Host FQDN**—Specify the fully qualified domain name (FQDN).

- **IP address, Gateway, Netmask, and Primary DNS**—Note that only IPv4 addresses are supported.
- **Network**—Displays all the networks that are available under the selected ESXi host or cluster.
For virtual networks (VLANs), this network carries Management traffic.
- **Data Store**—Displays all datastores that are accessible to the selected ESXi host or cluster. Select the datastore.

4. Click **Deploy**.

Results

PowerProtect Data Manager starts the configuration process. Go to **Reporting Engine** to check the status. You can also go to the **System Jobs** window to monitor the progress of the configuration job.

When the process is complete, a notification appears in the **Reporting Engine** window to indicate that the configuration is successful. You can now access reports from **Reports > Report Browser**.

Updating the reporting engine from version 19.10

Unless certain procedures are followed when updating PowerProtect Data Manager from version 19.10, any deployed reporting engine fails to update.

If you have deployed the reporting engine and are updating from PowerProtect Data Manager version 19.10, you must decide if you want to keep existing reporting data or delete it.

Keep reporting data

To update PowerProtect Data Manager and keep reporting data, see KB article 000199837: *PowerProtect Data Manager (PPDM) 19.10 Reporting update procedure*.

Delete reporting data

To update PowerProtect Data Manager and delete reporting data, follow these steps:

1. Delete the reporting engine. For more information, see *Deleting the reporting engine*.
2. Install the PowerProtect Data Manager update package.
3. Reconfigure and redeploy the reporting engine.

Report Browser

Use the **Report Browser** to view detailed reports for the data protection activities in your environment.

Report templates

Report templates are used to generate reports. When a template is selected, a particular report type is used. This report type is further modified by applying filters. Report templates and reports belong to either a job-activity category or an asset-protection category.

Built-in report templates

The **Report Browser** provides six job-activity built-in report templates and four asset-protection built-in report templates. If no reports are listed in a tab in the **Report Browser** pane, the built-in report templates are displayed. If one or more reports are listed in a tab in the **Report Browser** pane, the built-in report templates are displayed when a report is generated.

Built-in report templates cannot be edited or deleted, but reports and custom report templates based on them can be edited or deleted.

Custom report templates

Custom report templates can be created, edited, and deleted:

- When a report is generated, a custom report template associated with it is also created.
- If you edit a report from the **Report Browser** pane, the associated report template is updated.
- To view all custom report templates, go to **Reports > Report Templates**.
- To edit a custom report template:
 1. Click the name of the template from the **Report Templates** pane.
 2. After selecting the template, the UI changes to the **Report Browser** pane with the tab for the associated report selected.
 3. Click  and select **Edit**.
 4. To change the name of the report, click  to the right of the name, edit the name, and then click .
 5. To change the description of the report, click  to the right of the description, edit the description, and then click .
 6. Update the filters to apply to the report and click **Apply**.
 **NOTE:** To reset the filters to their default value, click **Reset**.
 7. Click  and select **Save Template**.
- To delete a custom report template, select the radio button to the left of its entry in the **Report Templates** pane and click **Delete**.

Reports

Learn about the reports that are available in the **Report Browser**.

The following figure provides an example **Jobs Status Summary** report.



Figure 15. Jobs Status Summary report






For each report, you can:

- Filter reports by choosing specific metrics.
- Show detailed information of a summary report.

Generate a report






To generate a report, perform the following actions:

1. To generate a report based on a built-in report template:

- a. Go to **Reports > Report Browser**
 - b. Click **Generate Report** under the template that the report should be based on.
- ① **NOTE:** If you do not see the built-in report templates, click 
2. To generate a report based on a custom report template:
 - a. Go to **Reports > Report Templates**.
 - b. Click the name of the template that the report should be based on.
 3. To change the name of the report, click  to the right of the name, edit the name, and then click .
 4. To change the description of the report, click  to the right of the description, edit the description, and then click .
 5. Select the filters to apply to the report and click **Apply**.
- ① **NOTE:** To reset the filters to their default value, click **Reset**.

Edit a report

To edit a report:

- From the **Report Browser** pane, select the tab of the report.
 - Click  and select **Edit**.
 - To change the name of the report, click  to the right of the name, edit the name, and then click .
 - To change the description of the report, click  to the right of the description, edit the description, and then click .
 - Update the filters to apply to the report and click **Apply**.
- ① **NOTE:** To reset the filters to their default value, click **Reset**.

Data collection frequency

PowerProtect Data Manager collects report data at regular intervals. The following table provides information about the type of data that PowerProtect Data Manager collects and the data collection frequency.

Table 44. Data collection frequency

Type of data	Description	Data collection frequency
Status	Overall status of the PowerProtect Data Manager server.	Every 15 minutes.
Configuration	Information about assets.	Every hour.
Protection jobs	Information about data protection activities, including Protect , Restore , and Replicate jobs.	Every 5 minutes.

- ① **NOTE:** Report data is not live and is as up-to-date as the last successful data collection request. Therefore, reports should be used for historical purposes only.
- To view live jobs data, go to **Jobs > Protection Jobs**.
 - To view live asset data, go to **Infrastructure > Assets**.
 - For a high-level view of the overall state of the PowerProtect Data Manager system, go to **Dashboard**.

Detailed report information and report timing

When you view detailed information by clicking the chart in a summary report, a new report is run and the latest data is displayed. Depending on when the two reports are run, this can result in the detailed report providing a job count that is different from the job count provided by the summary report. If the job counts are different, the job count provided by the detailed report is accurate. You can refresh the summary report page to update its information.

Report Browser options

From the **Report Browser** pane, click  to configure options for your reports.

The following table describes the menu items for reports:


Table 45. Report options

Menu item	Select the menu item to:
Edit	Configure filters and customization options.
Email	Email the report to one or more recipients.
Export	Export the report to a .csv file.
Save Template	Save the template of this report.
Schedule	Automatically send the report to one or more email recipients on a recurring schedule.

Emailing a report to one or more recipients


To send a report as a .csv file attachment, perform the following steps:


 **NOTE:** SMTP must be configured before performing these actions. For more information, see [Set up the email server](#).

1. From the **Report Browser** pane, select the tab of the report.
2. Click  and select **Email**. The **Email Report** window opens.
3. Provide information for **Report Name**.
4. Provide information for **To**, **Subject**, and **Body**.
5. Click **Send**.

Scheduling automatic reporting

To automatically send a report as a .csv file attachment on a recurring schedule, perform the following steps:

1. To send a report based on a custom report template, from the **Report Templates** pane, click the name of the template that the report should be based on.
2. From the **Report Browser** pane, select the tab of the report.
3. Click  and select **Schedule**. The **Schedule Report** pane opens.
4. Provide information for **Report Name** and **Schedule Name**.
5. Select a daily, weekly, or monthly schedule from the **Frequency** drop-down list.
6. Provide information for the time of day from the **At** drop-down lists.
7. If a weekly or monthly schedule is selected, select the day of the week or month from the appropriate controls.
8. Click **Next**.
9. Provide information for **To**, **Subject**, and **Body**.
10. Click **Next** to view a summary of the report name, schedule, and email details.
11. Click **Set Schedule**.

 **NOTE:** After automatic reporting has been scheduled for a report, you can see information about the report and its schedule by selecting the tab with its name from the **Report Browser** pane or its associated template in the **Report Templates** pane. When viewing the schedule information, you can disable or enable the schedule, delete the schedule, or edit the schedule.

Types of reporting information

Different reports provide different kinds of information. It is useful to know what types of information are available.

The following tables describe the different types of reports and what information they provide. The tables are grouped by category.

NOTE: This information is available across multiple reports and configurations. Individual reports might contain a subset of the information.

Table 46. Report types by job activity

Report type	Information displayed
Job Status Summary	The total number of successful and failed backup, restore, or replication jobs, along with a percentage summary.
Job Status by Asset Type	The total number of successful and failed backup, restores, or replication jobs, based on asset type.
Time-based Job Status	The number of successful and failed backup, restore, or replication jobs, restore jobs over a period of time.
Data Transfer Rate	The rate of data transfer over a period of time.
Asset Failure Rate	The assets with the highest count of consecutive primary backup failures, indicating the number of failures along with the time of the last successful backup or restore.
Jobs Summary - Table View	The details and status of all jobs, including: <ul style="list-style-type: none">• Asset Name• Asset Type• Host• Start Time• Job Status• Policy Name• Data Transferred

Table 47. Report types by asset protection

Report type	Information displayed
Asset Protection	A summary of the total number of protected and unprotected assets, as well as those in an Exclusion protection policy.
Asset Protection by Asset Type	A summary of the total number of protected and unprotected assets, as well as those in an Exclusion protection policy. The information is grouped by asset type.
Time-based Asset Protection	A summary of the last 7 days of the total number of protected and unprotected assets, as well as those in an Exclusion protection policy. The information is graphed.
Asset Jobs Distribution	The details and status of all jobs, including: <ul style="list-style-type: none">• Name• Asset Type• Host• Policy Name• Self Service• Last Copy• Asset Status• Protection Status

Filtering and customizing reports

The **Report Browser** provides options to filter and customize report data.

Filters that are applied to open reports are retained for the duration of the browser session. However, if a report is closed and then reopened during the same browser session, applied filters are not retained.

Using the search function when adding assets to a custom scope

You can use the search function to select assets in a custom scope. If you perform a search, only the assets selected from the search results are added to the custom scope.

CAUTION: If you edit an existing custom scope, using the search function removes the previously selected assets. To add new assets and keep the previously selected assets, do not use the search function. Add new assets by scrolling through the list of all assets and changing the current selections.

Deleting the reporting engine

Review the following information about deleting the reporting engine.

CAUTION: Deleting the reporting engine deletes all report data.

It is recommended that you do not delete the reporting engine.

To delete the reporting engine from PowerProtect Data Manager, go to **Reports > Reporting Engine** and click **Delete**. A notification appears in the window to indicate that deleting the reporting engine results in data loss.

Reconfiguring the reporting engine after deletion

To reconfigure the reporting engine, go to **Reports > Reporting Engine** and click **Configure**. For detailed steps on how to configure the reporting engine, go to **Configure** and deploy the reporting engine.

Managing disaster recovery of the reporting engine

As an administrator, you want to ensure that the reporting engine is protected from a disaster.

NOTE: Only the DD Boost storage type supports reporting server disaster recovery (DR). NFS is not supported.

When the reporting engine is deployed, the following occurs:

- The reporting engine and all reporting data are automatically backed up with configured server DR backups.
- If PowerProtect Data Manager is recovered from a server DR backup, the reporting engine and all reporting data are also recovered.

Recover the reporting engine from a DR backup


PowerProtect Data Manager automatically restores the reporting engine after disaster recovery of the PowerProtect Data Manager system is complete. If the PowerProtect Data Manager system could not restore the reporting engine automatically, use the steps in this procedure to restore only the reporting engine through the REST API. Recovery of a reporting engine must be performed on an operational PowerProtect Data Manager system. Only the Administrator role can restore the reporting engine.

Prerequisites

Obtain the name of the reporting engine backup from **System Settings > Disaster Recovery > Manage Backups**.

About this task

Use the backup manifest file to create a new text document that will be used issue a POST command with the REST API:

 CAUTION: Do not edit the manifest file itself.

Steps

1. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.
Use the same credentials that you used before PowerProtect Data Manager was restored.
2. Connect to the PowerProtect Data Manager console as an admin user.
3. Change directories to `/data01/server_backups/<PowerProtect Data Manager Hostname>_<NodeID>` to locate the backup manifest file.


Normally, there is only a single subdirectory in `/data01/server_backups`, so change to that subdirectory. However, if there is more than one subdirectory and you don't know which `<NodeID>` is the correct one, perform the following substeps:

- a. From `/data01/server_backups`, run the following commands, changing the username and password as required:

```
TOKEN=$(curl -X POST https://localhost:8443/api/v2/login -k -d '{ "username":  
"admin", "password": "admin_password" }' --header "Content-Type: application/json" |  
python3 -c "import sys, json; print(json.load(sys.stdin)['access_token'])")
```

```
curl -X GET https://localhost:8443/api/v2/nodes -k --header "Content-Type:  
application/json" --header "Authorization: Bearer $TOKEN"
```

- b. Run the command `grep -Rnwa -e '<Name>' --include=*.manifest`.
4. Copy the manifest file to a temporary file.
 5. Open the temporary file.
 6. Review the following example, and make the changes documented by the `//` comment entries.

 **NOTE:** The `//` comment entries displayed here do not exist in the temporary file itself. These comment entries are displayed here only as a guide.

```
{  
  "id": "ca8cbb13-6f3d-4ac5-87e5-de47a634379f",  
  "jobId": "990b4ea7-c0e4-4069-8dd5-7d0e084370fc", // DELETE LINE  
  "creationTime": "34e1c9dd-1b54-48b4-8283-151331d193ff",  
  "lastUpdated": "2022-08-25T19:40:18.165497Z", // DELETE LINE  
  "elapsedSeconds": 115,  
  "sequenceNumber": 89  
  "state": "Successful", // DELETE LINE  
  "version": "19.12.0-1-SNAPSHOT", // DELETE LINE  
  "hostname": "ldpdb141.hop.lab.emc.com", // DELETE LINE  
  "name": "mercijtestDr", // DELETE LINE  
  "nodeId": "a8d2df8e-5c3e-4160-87d4-32b9bfe6c283", // DELETE LINE  
  "sizeInBytes": 18244130,  
  "consistency": "CRASH CONSISTENT", // DELETE LINE  
  "checksum": "bbd97a04f296a8ed116e4a9272982d9e8411f3d0cf50dea131d5c2cd4ce224f8", //  
  DELETE LINE  
  "backupConsistencyType": "FULL", // DELETE LINE  
  "esSnapshotState": "UNKNOWN", // DELETE LINE  
  "backupTriggerSource": "USER", // DELETE LINE  
  "configType": "standalone", // DELETE LINE  
  "deployedPlatform": "vmware", // DELETE LINE  
  "replicationTargets": [], // DELETE LINE  
  "repositoryFileSystem": "BCOST FILE_SYSTEM", // DELETE LINE  
  "ddHostname": "ldpdg251.hop.lab.emc.com", // DELETE LINE and add line  
  "recover": true,  
  "Components": [ // change Components to components with lower case c  
    | // DELETE WHOLE PPDm COMPONENT LEAVING ONLY REPORTING  
    {  
      "name": "PPDM",  
      "id": "ca7cbb13-6f3d-4ac5-87e5-de47a634379f",  
      "lastActivityId": "2bdbe7a9-7c57-446d-b072-ad8081e2953d",  
      "version": "v2",  
      "backupPath": "ldpdg251.hop.lab.emc.com:SysDR_ldpdb141/  
ldpdb141_a8d2df8e-5c3e-4160-87d4-32b9bfe6c283/PPDM",  
      "backupStatus": "SUCCESSFUL",  
      "backupsEnabled": true,  
      "errorResults": []  
    }, // STOP DELETING HERE  
  ]  
}
```

```

    {
      "name": "REPORTING",
      "id": "34e1c9dd-1b54-48b4-8283-151331d193ff",
      "lastActivityId": "ed2dc805-clf7-42fd-b9af-71897fc1da01",
      "version": "v2",
      "backupPath": "192.168.100.109:SysDR_DPDII2201IDPA10/
ppdm_64d2f00a-lce0-47b5-9c60-914ea7d0ele8/REPORTING",
      "backupStatus": "SUCCESSFUL",
      "backupsEnabled": true, // DELETE TRAILING COMMA
      "errorResults": [] // DELETE LINE
    }
  ], // DELETE TRAILING COMMA
  "componentVersions": [], // DELETE LINE
  "expirationTime": "2023-06-11T09:41:20.383633Z", // DELETE LINE
  "protectionCopySetId": "07e7af37-1a80-5436-b320-9e537fba1317" // DELETE LINE
}

```

In summary:

- remove all lines with the // DELETE LINE comment entry displayed here
- add recover: true
- change Components to components
- remove all listed component blocks except for REPORTING
- remove the trailing comma from "backupsEnabled": true,
- remove the trailing comma from [,

The result of these changes should look similar to the following:

```

{
  "id": "ca8cbb13-6f3d-4ac5-87e5-de47a634379f",
  "creationTime": "2022-10-12T15:01:13.476401+0000",
  "elapsedSeconds": 115,
  "sequenceNumber": 89,
  "sizeInBytes": 18244130,
  "recover": true,
  "components": [
    {
      "name": "REPORTING",
      "id": "ca8cbb13-6f3d-4ac5-87e5-de47a634379f",
      "lastActivityId": "ed2dc805-clf7-42fd-b9af-71897fc1da01",
      "version": "v2",
      "backupPath": "192.168.100.109:SysDR_DPDII2201IDPA10/
ppdm_64d2f00a-lce0-47b5-9c60-914ea7d0ele8/REPORTING",
      "backupStatus": "SUCCESSFUL",
      "backupsEnabled": true
    }
  ]
}

```

7. Copy the value of the text inside the quotation marks that follow "id":.
This value replaces the variable <backupID> used in step 11. In this example, <backupID> is ca8cbb13-6f3d-4ac5-87e5-de47a634379f.
8. Remove all carriage returns from the temporary file, so that all the text is on a single line.
9. Copy all of the text from the temporary file.
This value replaces the variable <mon/feetText> used in step 11.
10. Run the following command, changing the username and password credentials as required:

NOTE: Even if you ran this command in step 3.a, run it again. The validity of the value of TOKEN is time sensitive.

```

TOKEN=$(curl -X POST https://localhost:8443/api/v2/login -k -d '{ "username":
"admin", "password": "admin_password" }' --header "Content-Type: application/json" |
python3 -c "import sys, json; print(json.load(sys.stdin)['access_token'])")

```

11. Run the following command:

```
curl -X PUT 'https://localhost:8443/api/v2/server-disaster-recovery-backups/  
<backupID>' --header "Authorization: Bearer $TOKEN" --header 'Content-Type:  
application/json' -k -d '<manifestText>'
```

- Replace *<backupID>* with the value obtained in step 7.
- Replace *<manifestText>* with all of the text obtained in step 9.

12. To monitor the status of the restore process in the PowerProtect Data Manager UI, select **Jobs > System Jobs** and look for a job with the description **Server Disaster Recovery Restore**.

Next steps

Delete the temporary file created in step 4.

Configuring and Managing the PowerProtect Agent Service

Topics:

- About the PowerProtect agent service
- Start, stop, or obtain the status of the PowerProtect agent service
- Register the PowerProtect agent service to a different server address
- Recovering the PowerProtect agent service from a disaster
- Troubleshooting agent registration

About the PowerProtect agent service

The PowerProtect agent service is a REST API based service that is installed by the application agent on the application host. The agent service provides services and APIs for discovery, protection, restore, instant access, and other related operations. The PowerProtect Data Manager uses the agent service to provide integrated data protection for the application assets.

This section uses `<agent_service_installation_location>` to represent the PowerProtect agent service installation directory. By default, the agent service installation location is `C:\Program Files\DPSAPPS\AgentService` on Windows and `/opt/dpsapps/agentsvc` on Linux. All files that are referenced in this section are the relative paths to the agent service installation location.

The PowerProtect agent service performs the following operations:

- **Addon detection**—An addon integrates the application agent into the agent service. The agent service automatically detects the addons on the system for each application asset type and notifies the PowerProtect Data Manager. While multiple addons can operate with different asset types, only one agent service runs on the application host. Specific asset types can coexist on the same application host.
- **Discovery**—The agent service discovers both stand-alone and clustered database servers (application systems), databases and file systems (assets), and their backup copies on the application agent host. After the initial discovery, when the agent service discovers any new application systems, assets, or copies, the agent service notifies the PowerProtect Data Manager.
- **Self-service configuration**—The agent service can configure the application agent for self-service operations by using information that is provided by the PowerProtect Data Manager. When you add an asset to a protection policy for self-service or centralized protection, or modify the protection policy, including changing the DD Boost credentials, the PowerProtect Data Manager automatically pushes the protection configuration to the agents.
 - ⓘ **NOTE:** If you change the DD Boost credentials to include `\` in the password, the protection policy configuration will not be pushed to the agents unless you also select the protection policy from the **Protection Policies** window, and then click **Set LockBox**.
- **Centralized backups**—The agent service performs the centralized backups as requested by the PowerProtect Data Manager.
- **Centralized restores**—The agent service performs the centralized restores as requested by the PowerProtect Data Manager.
 - ⓘ **NOTE:** In the current release, the centralized restores are only available for the File System agent, Microsoft SQL Server agent, and Storage Direct agent.
- **Backup deletion and catalog cleanup**—The PowerProtect Data Manager deletes the backup files directly from the protection storage when a backup expires or an explicit delete request is received and no dependent (incremental or log) backups exist. The PowerProtect Data Manager goes through the agent service to delete the catalog entries from the database vendor's catalog and the agent's local datastore.
 - ⓘ **NOTE:** Deletion of any backup copies manually or through the command line is not recommended. PowerProtect Data Manager deletes all the expired copies as needed.

The agent service maintains SQLite database backups in the `<install_directory>/dbs/v1/backups` directory, which is cleaned based on the retention time in the `config.yml` file. The agent service cleans up the backups only when the backup count exceeds 10 (cleans up only extra backups after the 10th count).

The agent service is started during the agent installation by the installer. The agent service runs in the background as a service and you do not interact with it directly.

The `config.yml` file contains the configuration information for the agent service, including several parameter settings that you can change within the file. The `config.yml` file is located in the `<agent_service_installation_location>` directory.

If the `config.yml` file becomes corrupted, you can run the following commands to restore the file and continue the protection provided by the agent service:

- On Windows:

```
agentService.exe config=config.yml service=false restoreConfig=true
```

- On Linux and AIX:

```
agentService config=config.yml service=false restoreConfig=true
```

The agent service periodically starts subprocesses to perform the discovery jobs. You can see the type and frequency of these jobs in the `jobs:` section of the `config.yml` file. The job interval unit is minutes.

The agent service maintains a datastore in the `<agent_service_installation_location>/dbs/v1` directory, which contains information about the application system, assets, and backups discovered on the system. The size of the datastore files depends on the number of applications and copies on the host. The agent service periodically creates a backup of its datastore in the `<agent_service_installation_location>/dbs/v1/backups` directory, as used to recover the datastore if this datastore is lost.

NOTE: The size of each datastore backup is the same as the datastore itself. By default, a backup is created every hour. To save space on the file system, you can reduce this datastore backup frequency for large datastores. By default, the datastore backup is retained for one week. You can change the datastore backup frequency, retention period, and backup location in the `config.yml` file.

Start, stop, or obtain the status of the PowerProtect agent service

The PowerProtect agent service is started during the agent installation by the installer. If needed, you can use the appropriate procedure to start, stop, or obtain the status of the agent service.

On Linux, you can start, stop, or obtain the status of the agent service by running the `register.sh` script that is found in the `<agent_service_installation_location>` directory.

- To start the agent service:

```
# register.sh --start
Started agent service with PID - 1234
```

- To stop the agent service:

```
# register.sh --stop
Successfully stopped agent-service.
```

- To obtain the status when the agent service is running:

```
# register.sh --status
Agent-service is running with PID - 1234
```

- To obtain the status when the agent service is not running:

```
# register.sh --status  
  
Agent-service is not running.
```

On Windows, you can start, stop, or obtain the status of the PowerProtect agent service from the Services Manager, similar to other Windows services. The name of the service in Services Manager is **PowerProtect Agent Service**.

Register the PowerProtect agent service to a different server address

The PowerProtect agent service is registered to a particular PowerProtect Data Manager server during the agent installation by the installer. If needed, you can register the agent service to a different PowerProtect Data Manager server address.

The agent service can only be registered to a single PowerProtect Data Manager server. When you register the agent service to a new server, the agent service will automatically unregister from the previous server address.

Before you register the agent service to a new server, ensure that you complete the following steps:

1. Stop the agent service as described in the preceding topic.
2. Delete the `<agent_service_installation_location>/ssl` folder and `<agent_service_installation_location>/dbs/v1/objects.db`.

On Linux, you can register the agent service to a different server address by running the `register.sh` script that is found in the `<agent_service_installation_location>` directory.

NOTE: The `register.sh` script stops the currently running agent service.

- The following command prompts for the new IP address or hostname:

```
# register.sh  
  
Enter the PowerProtect Data Manager IP address or hostname: 10.0.01  
  
Warning: Changing IP of PowerProtect Server from 192.168.0.1 to 10.0.0.1  
  
Started agent service with PID - 1234
```

- The following command includes the new IP address on the command line:

```
# register.sh --ppdmServer=10.0.0.1  
  
Warning: Changing IP of PowerProtect Server from 192.168.0.1 to 10.0.0.1  
  
Started agent service with PID - 1234
```

On Windows, you can change the PowerProtect Data Manager server address by launching the agent installer and selecting the change option. Change the PowerProtect Data Manager service address from the **Configuration Install Options** page.

Recovering the PowerProtect agent service from a disaster

You can perform self-service restores of application assets by using a file system or application agent, regardless of the state of the agent service or PowerProtect Data Manager. The information in this section describes how to bring the agent service to an operational state to continue if a disaster occurs and the agent service datastore is lost.

The agent service periodically creates a backup of its datastore in the `<agent_service_installation_location>/dbs/v1/backups` repository. If all of these backups are lost, the agent service can still start. The agent service discovers all the application systems, assets, and backup copies on the system again, and notifies PowerProtect Data Manager. Depending on when the failure occurred, the agent service might not be able to find older backup copies for some asset types. As a result, the centralized deletion operations might fail when cleaning up the database vendor catalog or removing older backups that are taken before the asset is added to PowerProtect Data Manager.

By default, the agent service backs up consistent copies of its datastore files to the local disk every hour and keeps the copies for 7 days. Each time the agent service backs up the contents of the datastore, it creates a subdirectory under the `<agent_service_installation_location>/dbs/v1/backups` repository. The subdirectories are named after the time the operation occurred, in the format `YYYY-MM-DD_HH-MM-SS_epochTime`.

By default, the datastore repository is on the local disk. To ensure that the agent service datastore and its local backups are not lost, it is recommended that you back up the datastore through file system backups. You can also change the datastore backup location to a different location that is not local to the system. To change the datastore backup location, update the values in the `config.yml` file.

Restore the PowerProtect Data Manager agent service datastore

Prerequisites

NOTE: Ensure that the agent service is powered off. Do not start the agent service until disaster recovery is complete.

About this task

You can restore the datastore from the datastore backup repository. If the repository is no longer on the local disk, restore the datastore from file system backups first.

To restore the datastore from a backup in the datastore backup repository, complete the following steps:

Steps

1. Move the files in the `<agent_service_installation_location>/dbs/v1` directory to a location for safe keeping.

NOTE: Do not move or delete any `<agent_service_installation_location>/dbs/v1` subdirectories.

2. Select the most recent datastore backup.

The directories in the datastore backup repository are named after the time the backup was created.

3. Copy the contents of the datastore backup directory to the `<agent_service_installation_location>/dbs/v1` directory.

After the copy operation is complete, the `<agent_service_installation_location>/dbs/v1` directory should contain the following files:

- `copies.db`
- `objects.db`
- `resources.db`
- `sessions.db`

4. Start the agent service.

Troubleshooting agent registration

Review the following information that is related to troubleshooting agent registration issues.

On Windows, if the agent fails to establish a connection with the PowerProtect Data Manager server, agent registration might fail with the following error message:

```
During a network connectivity test, the agent is unable to reach the PowerProtect Data Manager server by using ping.
```

1. If the ping command is blocked in the environment, the agent registration can still complete successfully. Review the agent service logs at `INSTALL_DIR\DPSAPPS\AgentService\logs` to verify that the registration is successful. If the registration is successful, the status of the agent host indicates **Registered** in the PowerProtect Data Manager UI.
2. If the ping command is not blocked in the environment, the agent registration might not complete successfully because a network connection cannot be started. If this occurs, complete the following steps to troubleshoot the issue:

On Linux or AIX, if the agent fails to establish a connection with the PowerProtect Data Manager server, agent registration might fail with the following error message:

During a network connectivity test, the agent is unable to reach the PowerProtect Data Manager server by using ping and curl.

1. If the ping command is blocked in the environment and curl is not installed, the agent registration can still complete successfully. Review the agent service logs at `/opt/dpsepps/agentsvc/logs` to verify that the registration is successful. If the registration is successful, the status of the agent host indicates **Registered** in the PowerProtect Data Manager UI.
2. If the ping command is not blocked in the environment, the agent registration might not complete successfully because a network connection cannot be started. If this occurs, complete the following steps to troubleshoot the issue:

If agent registration fails with these error messages, complete the following operation:

1. Use any network packet tracing tool to trace the packets from the agent system to PowerProtect Data Manager.
2. Start the packet tracing between the source IP of the agent system and the destination IP of PowerProtect Data Manager.
3. Start the network traffic between the agent system and PowerProtect Data Manager.

Wait 10 to 15 seconds.

4. Analyze the captured packets.
5. Look for SYN and SYN_ACK packets to see if a 3-way handshake is being performed.

Determine whether the source agent or the destination PowerProtect Data Manager is blocking the connection.

If network traffic is blocked, contact your network security team to resolve the port communication issue.

Glossary

This glossary provides definitions of acronyms used across the product documentation set:

A

AAG: Always On availability group

ACL: access control list

AD: Active Directory

AKS: Azure Kubernetes Service

API: application programming interface

ARM: Azure Resource Manager

AVS: Azure VMware Solution

AWS: Amazon Web Services

AZ: availability zone

B

BBB: block-based backup

C

CA: certificate authority

CBT: Changed Block Tracking

CDC: change data capture

CIFS: Common Internet File System

CLI: command-line interface

CLR: Common Language Runtime

CN: common name

CPU: central processing unit

CR: custom resource

CRD: custom resource definition

CSI: container storage interface

CSV: Cluster Shared Volume

D

DAG: database availability group

DBA: database administrator

DBID: database identifier

DDMC: DD Management Center

DDOS: DD Operating System

DDVE: DD Virtual Edition

deploy

At Dell, virtual machines are deployed to virtual environments, while software components and hardware devices are installed. Both PowerProtect Data Manager and DDVE are virtual machines that are deployed. If you are searching this software guide for instances of `install` and not finding anything appropriate, search for `deploy` instead.

DFC: DD Boost over Fibre Channel

DNS: Domain Name System

DPC: Data Protection Central

DR: disaster recovery

DRS: Distributed Resource Scheduler

DSA: Dell security advisory

E

EBS: Elastic Block Store

EC2: Elastic Compute Cloud

eCDM: Enterprise Copy Data Management

ECS: Elastic Cloud Storage

EFI: Extensible Firmware Interface

EKS: Elastic Kubernetes Service

ENI: Elastic Network Interface

EULA: end-user license agreement

F

FC: Fibre Channel

FCD: first class disk

FCI: failover cluster instance

FETB: front-end protected capacity by terabyte

FLR: file-level restore

FGDN: fully qualified domain name

FTP: File Transfer Protocol

G

GB: gigabyte

At Dell, this is 2^{30} bytes.

Gb/s: gigabits per second

At Dell, this is 2^{30} bits per second.

GCP: Google Cloud Platform

GCVE: Google Cloud Virtual Edition

GID: group identifier

GLR: granular-level restore

GUI: graphical user interface

GUID: globally unique identifier

H

HA: High Availability

HANA: high-performance analytic appliance

HTML: Hypertext Markup Language

HTTP: Hypertext Transfer Protocol

HTTPS: Hypertext Transfer Protocol Secure

I

IAM: identity and access management

IDE: Integrated Device Electronics

IP: Internet Protocol

IPv4: Internet Protocol version 4

IPv6: Internet Protocol version 6

K

KB: kilobyte

At Dell, this is 2^{10} bytes.

L

LAC: License Authorization Code

LAN: local area network

M

MB: megabyte

At Dell, this is 2^{20} bytes.

ns: millisecond

MTU: maximum transmission unit

N

NAS: network-attached storage

NBD: network block device

NBDSSL: network block device over SSL

NDMP: Network Data Management Protocol

NFC: Network File Copy

NFS: Network File System

NIC: network interface card

NTFS: New Technology File System

NTP: Network Time Protocol

O

OS: operating system

OSS: open-source software

OVA: Open Virtualization Appliance

P

PCS: Protection Copy Set

PDF: Portable Document Format

PEM: Privacy-enhanced Electronic Mail

PIN: personal identification number

PIT: point in time

PKCS: Public Key Cryptography Standards

PSC: Platform Service Controller

PVC (cloud computing): private virtual cloud

PVC (Kubernetes): Persistent Volume Claim

R

RAC: Real Application Cluster

RAM: random-access memory

RBAC: role-based access control

ReFS: Resilient File System

REST API: representational-state transfer API

RHEL: RedHat Enterprise Linux

RMAN: Recovery Manager

RPO: recovery-point objective

RSA: Rivest-Shamir-Adleman

S

S3: Simple Storage Services

SaaS: software as a service

SAP: System Analysis Program Development
From the SAP website (2022), "the name is an initialism of the company's original German name: Systemanalyse Programmentwicklung, which translates to System Analysis Program Development. Today the company's legal corporate name is SAP SE - SE stands for *societas Europaea*, a public company registered in accordance with the European Union corporate law."¹

SCSI: Small Computer System Interface

SDDC: software-defined data center

SELinux: Security-Enhanced Linux

SFTP: Secure File Transfer Protocol

SLA: service-level agreement

SLES: SUSE Linux Enterprise Server

SLO: service-level objective

SPBM: Storage Policy Based Management

SQL: Structured Query Language

SRS: Secure Remote Services

SSD: solid-state drive

SSH: Secure Shell

SSL: Secure Sockets Layer

SSMS: SQL Server Management Studio

SSVs: System Stable Values

T

TB: terabyte
At Dell, this is 2^{40} bytes.

TCP: Transmission Control Protocol

TDE: Transparent Data Encryption

TLS: Transport Layer Security
TPM: Trusted Platform Module
TSDM: Transparent Snapshots Data Mover
T-SQL: Transact-SQL

U

UAC: user account control
UDP: User Datagram Protocol
UI: user interface
UID: user identifier

update

At Dell, software is updated and hardware is upgraded. If you are searching this software guide for instances of upgrade and not finding any, search for update instead.

UTC: Coordinated Universal Time

From Wikipedia (2022), "this abbreviation comes as a result of the International Telecommunication Union and the International Astronomical Union wanting to use the same abbreviation in all languages. English speakers originally proposed CUT (for 'coordinated universal time'), while French speakers proposed TUC (for 'temps universel coordonné')."

V

VADP: VMware vStorage API for Data Protection
VBS: virtualization-based security
VCF: VMware Cloud Foundation
vCLS: vSphere Cluster Service
vCSA: vCenter Server Appliance
VCSA: vCenter Server Appliance
VDI: Virtual Device Interface
vDisk: virtual disk
vDS: virtual distributed switch
vFRC: Virtual Flash Read Cache
VGT: Virtual Guest Tagging
VIB: vSphere Installation Bundle
VLAN: virtual LAN
VM: virtual machine
VMC: VMware Cloud
VMDK: virtual machine disk
VNet: virtual network

VPC: virtual private cloud
VRSLCM: vRealize Suite Lifecycle Manager
VST: Virtual Switch Tagging
vTPM: Virtual Trusted Platform Module
VVD: VMware Validated Design
vVol: virtual volume

W

WAN: wide area network

Item 2 – PPDM (Cont...)

DECISION

Brasil (Sede)

Sede Financeira S/A - Quadra 06 - Conjunto W
Bloco A - Sala 102 - Asa B2 - Brasília/DF
Cap. 70.329-915 - Tel. (61) 3045-0050

Salvador

Avenida Tancredo Neves, 626 - Salas 2010 e 2011
2º andar - Torre Empresarial do Ecl. Mundo Plaza
Camino das Ávoas - Salvador/BA - Cap. 41.800-020
Tel. (71) 3565.7007

São Paulo

Rua Arzuna, 1.422 - Conjunto 75 - Ed. Platium
Building Semir - Berrini - São Paulo/SP - Cap. 04.567-005
Tel. (11) 5583.0544



PowerProtect Data Manager 19.14

File System User Guide

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Contents

Preface	6
Chapter 1: PowerProtect Data Manager File System Agent Overview	11
PowerProtect Data Manager overview.....	11
Introducing the File System agent.....	11
File System agent system requirements.....	12
Supported Internet Protocol versions.....	12
Encryption in-flight.....	15
PowerProtect Data Manager new deployment overview.....	15
PowerProtect Data Manager existing deployment overview.....	16
Security configuration.....	16
Role-based security.....	17
Chapter 2: Enabling the File System Agent	18
About the File System agent.....	18
Application agent and File System agent coexistence.....	18
File System agent prerequisites.....	20
File System agent limitations.....	21
Block-based backups.....	22
Best practices for file system backups.....	22
Configure the file system parallel backup setting.....	23
Configure asset multi-streaming for file-based backups.....	24
Protect an asset with the File System agent.....	24
Protect an asset in a Microsoft Windows Server clustered environment with the File System agent.....	25
Installing and uninstalling the File System agent on AIX.....	25
Install the File System agent on AIX.....	25
Uninstall the File System agent on AIX.....	27
Installing and updating the File System agent on Linux.....	28
Install the File System agent on Linux.....	28
Update the File System agent on Linux.....	30
Uninstall the File System agent on Linux.....	31
Reregister the File System agent on Linux or AIX.....	31
Installing and updating the File System agent on Windows.....	32
Install the File System agent on Windows.....	32
Update the File System agent on Windows.....	34
Uninstall the File System agent on Windows.....	36
Reregister the File System agent on Windows.....	37
Manage the File System agent.....	37
Changing the preferred host address.....	38
View application agent details.....	39
Configure the centralized file-level restore service port number.....	40
Update the application agent in the PowerProtect Data Manager UI.....	41
Moving a File System agent from one PowerProtect Data Manager system to another.....	42
.....	42

Enable the File System agent after migrating File System assets to a refreshed server.....	43
Chapter 3: Managing Storage, Assets, and Protection.....	44
Enable an asset source.....	44
Disable an asset source.....	45
Delete an asset source.....	45
Discover a file system host.....	46
Protection policies for File System agent.....	46
Supported protection policy purposes.....	46
Supported protection policy objectives.....	47
File System replication triggers.....	48
Roadmap for planning a protection policy.....	48
Before you add a protection policy for file system protection.....	49
Add a protection policy for file system protection.....	53
Protection rules.....	61
Protection rule attributes and criteria.....	62
Add a protection rule.....	64
Manually run a protection rule.....	66
Edit or delete a protection rule.....	66
View assets applied to a protection rule.....	67
Change the priority of an existing protection rule.....	67
Configure protection rule behavior.....	67
Cancel a File System agent backup or restore job.....	68
Add a service-level agreement.....	69
Extended retention for protection policies created in PowerProtect Data Manager 19.11 or earlier.....	71
Edit the retention period for backup copies.....	73
Delete backup copies.....	74
Retry a failed backup copy deletion.....	75
Export data for deleted Exchange, File System, Kubernetes, Block Volume, and SQL backup copies.....	75
Remove Exchange, File System, Kubernetes, Block Volume, and SQL backup copies from the PowerProtect Data Manager database.....	76
Host CPU throttling.....	76
Exclusion filters for File Systems.....	78
Add an exclusion filter.....	78
Edit or delete an exclusion filter.....	81
Apply an exclusion filter to a protection policy.....	81
Remove an exclusion filter from a protection policy.....	82
Centralized restore of a file-system asset.....	82
Centralized image-level restore of a file-system asset.....	83
Centralized file-level restore of a file-system asset.....	84
Updating the File System agent hostname or IP address.....	88
Update the File System agent hostname.....	88
Update the File System agent IP address, or reregister the agent after an agent host or PowerProtect Data Manager IP address change.....	88
Manage the PowerProtect agent service.....	89
About the PowerProtect agent service.....	89
Changing the preferred host address.....	90
Start, stop, or obtain the status of the PowerProtect agent service.....	90
Recovering the PowerProtect agent service from a disaster.....	91

Chapter 4: Performing Self-Service Backups and Restores with the File System Agent.....	93
Performing self-service backups of file systems.....	93
Performing self-service restore of a file system host.....	94
Using the ddfsadmin utility for file systems.....	94
Self-service image-level restore of file systems.....	94
Self-service file-level restore of file systems.....	95
Chapter 5: Performing Self-Service Backups and Restores for Disaster Recovery.....	97
Performing self-service backups for disaster recovery.....	97
Using the ddfsadmin utility for disaster recovery.....	98
Self-service system state restore for disaster recovery.....	98
Bare-metal recovery restore for disaster recovery.....	99
Chapter 6: Performing Disaster Recovery with the File System Agent in Windows.....	100
Disaster recovery limitations.....	100
Preparing for disaster recovery.....	100
Gathering key information.....	100
Critical volumes in disaster recovery	100
Discover the assets to back up.....	101
Create a disaster recovery protection policy.....	101
Synchronize all clocks.....	102
Manually run a disaster recovery policy.....	102
Performing system-state recovery.....	102
Perform a system-state recovery.....	102
Recovering the Active Directory.....	103
Configure the client to boot into Directory Services Restore Mode.....	103
Recover the Active Directory from disaster-recovery data.....	104
Authoritative and nonauthoritative Microsoft restores.....	104
Performing bare-metal recovery.....	105
Bare-metal recovery requirements.....	106
About the WinPE image.....	106
Using a custom WinPE image.....	106
Perform a bare-metal recovery.....	108
Perform bare-metal recovery of Windows clusters.....	111
Disaster recovery solution for Microsoft SQL application data.....	112
Performing DR for Microsoft SQL application data.....	112
Protecting Microsoft Distributed File System using BMR and SSR.....	113
Performing application restores after bare-metal recovery.....	114
Appendix A: File System Best Practices and Troubleshooting.....	115
Troubleshooting installation and operation.....	115
Troubleshooting backups.....	118
Troubleshooting disaster recovery.....	119
Troubleshooting restores.....	122
Troubleshooting storage units.....	124
Glossary.....	125

As part of an effort to improve product lines, periodic revisions of software and hardware are released. Therefore, all versions of the software or hardware currently in use might not support some functions that are described in this document. The product release notes provide the most up-to-date information on product features.

If a product does not function correctly or does not function as described in this document, contact Customer Support.

NOTE: This document was accurate at publication time. To ensure that you are using the latest version of this document, go to the Customer Support website.

Product naming

Data Domain (DD) is now PowerProtect DD. References to Data Domain or Data Domain systems in this documentation, in the user interface, and elsewhere in the product include PowerProtect DD systems and older Data Domain systems.

Isilon is now PowerScale. References to Isilon, Isilon products, or Isilon appliances in this documentation, in the user interface, and elsewhere in the product include PowerScale products and appliances.

In many cases the user interface has not yet been updated to reflect these changes.

Language use

This document might contain language that is not consistent with Dell Technologies current guidelines. Dell Technologies plans to update the document over subsequent future releases to revise the language accordingly.

This document might contain language from third-party content that is not under Dell Technologies control and is not consistent with the current guidelines for Dell Technologies own content. When such third-party content is updated by the relevant third parties, this document will be revised accordingly.

Acronyms

The acronyms used in this document might not be familiar to everyone. Although most acronyms are defined on their first use, a definition is not always provided with later uses of the acronym. For a list of all acronyms and their definitions, see the glossary at the end of the document.

Website links

The website links used in this document were valid at publication time. If you find a broken link, provide feedback on the document, and a Dell Technologies employee will update the link in the next release as necessary.

Purpose

This document describes how to configure and administer the Dell PowerProtect Data Manager to protect and recover data on the file system host. The *PowerProtect Data Manager Administration and User Guide* provides additional details about configuration and usage procedures.

Audience

This document is intended for the host system administrator who is involved in managing, protecting, and reusing data across the enterprise by deploying PowerProtect Data Manager.

Revision history

The following table presents the revision history of this document.

Table 1. Revision history

Revision	Date	Description
01	July 11, 2023	Initial release of this document for PowerProtect Data Manager version 19.14.

Compatibility information

Software compatibility information for the PowerProtect Data Manager software is provided by the E-Lab Navigator.

Related documentation

The following publications are available at Customer Support and provide additional information:

Table 2. Related documentation

Title	Content
<i>PowerProtect Data Manager Administrator Guide</i>	Describes how to configure, use and administer the software. This guide also includes disaster recovery procedures. Procedures specific to asset protection are provided in the individual user guides.
<i>PowerProtect Data Manager Deployment Guide</i>	Describes how to deploy and license the software.
<i>PowerProtect Data Manager Release Notes</i>	Contains information about new features, known limitations, environment, and system requirements for the software.
<i>PowerProtect Data Manager Security Configuration Guide</i>	Contains security information.
<i>PowerProtect Data Manager Amazon Web Services Deployment Guide</i>	Describes how to deploy the software to Amazon Web Services (AWS).
<i>PowerProtect Data Manager Azure Deployment Guide</i>	Describes how to deploy the software to Microsoft Azure.
<i>PowerProtect Data Manager Google Cloud Platform Deployment Guide</i>	Describes how to deploy the software to Google Cloud Platform (GCP).
<i>PowerProtect Data Manager Cloud Disaster Recovery Administration and User Guide</i>	Describes how to deploy Cloud Disaster Recovery (Cloud DR), protect virtual machines in the AWS or Azure cloud, and run recovery operations.
<i>PowerProtect Data Manager Cyber Recovery User Guide</i>	Describes how to install, update, patch, and uninstall the PowerProtect Cyber Recovery software.
<i>PowerProtect Data Manager File System User Guide</i>	Describes how to configure and use the software with the File System agent for file-system data protection.
<i>PowerProtect Data Manager Kubernetes User Guide</i>	Describes how to configure and use the software to back up and restore namespaces and PVCs in a Kubernetes cluster or Tanzu Kubernetes cluster.
<i>PowerProtect Data Manager Microsoft Exchange Server User Guide</i>	Describes how to configure and use the software to back up and restore the data in a Microsoft Exchange Server environment.
<i>PowerProtect Data Manager Microsoft SQL Server User Guide</i>	Describes how to configure and use the software to back up and restore the data in a Microsoft SQL Server environment.

Table 2. Related documentation (continued)

Title	Content
<i>PowerProtect Data Manager Oracle RMAN User Guide</i>	Describes how to configure and use the software to back up and restore the data in an Oracle Server environment.
<i>PowerProtect Data Manager SAP HANA User Guide</i>	Describes how to configure and use the software to back up and restore the data in an SAP HANA Server environment.
<i>PowerProtect Data Manager Storage Direct User Guide</i>	Describes how to configure and use the software with the Storage Direct agent to protect data on VMAX storage arrays through snapshot backup technology.
<i>PowerProtect Data Manager Network-Attached Storage User Guide</i>	Describes how to configure and use the software to protect and recover the data on network-attached storage (NAS) shares and appliances.
<i>PowerProtect Data Manager Virtual Machine User Guide</i>	Describes how to configure and use the software to back up and restore virtual machines and virtual machine disks (VMDKs) in a vCenter Server environment with VADP or the Transparent Snapshots Data Mover (TSDM).
<i>PowerProtect Data Manager Storage Array User Guide</i>	Describes how to configure and use the software to protect and restore data on PowerStore storage arrays.
<i>VMware Cloud Foundation Disaster Recovery With PowerProtect Data Manager</i>	Provides a detailed description of how to perform an end-to-end disaster recovery of a VMware Cloud Foundation (VCF) environment.
PowerProtect Data Manager Public REST API documentation	Contains the Dell Technologies APIs and includes tutorials to guide you in their use.
<i>vRealize Automation Data Protection Extension for Data Protection Systems Installation and Administration Guide</i>	Describes how to install, configure, and use the vRealize Data Protection Extension.

Typographical conventions

The following type style conventions are used in this document:

Table 3. Style conventions

Formatting	Description
Bold	Used for interface elements that a user specifically selects or clicks, for example, names of buttons, fields, tab names, and menu paths. Also used for the name of a dialog box, page, pane, screen area with title, table label, and window.
<i>Italic</i>	Used for full titles of publications that are referenced in text.
Monospace	Used for: <ul style="list-style-type: none"> • System code • System output, such as an error message or script • Pathnames, file names, file name extensions, prompts, and syntax • Commands and options
<i>Monospace italic</i>	Used for variables.
Monospace bold	Used for user input.
[]	Square brackets enclose optional values.
	Vertical line indicates alternate selections. The vertical line means or for the alternate selections.
{ }	Braces enclose content that the user must specify, such as x, y, or z.
...	Ellipses indicate non-essential information that is omitted from the example.

You can use the following resources to find more information about this product, obtain support, and provide feedback.

Where to find product documentation

To find the latest documentation, navigate to the PowerProtect Data Manager Info Hub or type `www.dell.com/ppdm/docs` in your browser, or scan the following QR code on your mobile device.



Where to get support

The Customer Support website provides access to product licensing, documentation, advisories, downloads, and how-to and troubleshooting information. The information can enable you to resolve a product issue before you contact Customer Support.

To access a product-specific page:

1. Go to the Customer Support website.
2. In the search box, type a product name, and then from the list that appears, select the product.

Support Library

The Support Library contains a knowledge base of applicable solutions that you can search for either by solution number (for example, KB000xxxxxx) or by keyword.

To search the Support Library:

1. Go to the Customer Support website.
2. On the **Support** tab, click **Support Library**.
3. In the search box, type either the solution number or keywords. Optionally, you can limit the search to specific products by typing a product name in the search box, and then selecting the product from the list that appears.

Live chat

To participate in a live interactive chat with a support agent:

1. Go to the Customer Support website.
2. On the **Support** tab, click **Contact Support**.
3. On the **Contact Information** page, click the relevant support, and then proceed.

Service requests

To obtain in-depth help from a support agent, submit a service request. To submit a service request:

1. Go to the Customer Support website.
 2. On the **Support** tab, click **Service Requests**.
- ① **NOTE:** To create a service request, you must have a valid support agreement. For details about either an account or obtaining a valid support agreement, contact a sales representative. To find the details of a service request, in the **Service Request Number** field, type the service request number, and then click the right arrow.

To review an open service request:

1. Go to the Customer Support website.
2. On the **Support** tab, click **Service Requests**.
3. On the **Service Requests** page, under **Manage Your Service Requests**, click **View All Dell Service Requests**.

Online communities

For peer contacts, conversations, and content on product support and solutions, go to the Community Network. Interactively engage with customers, partners, and certified professionals online.

How to provide feedback

Feedback helps to improve the accuracy, organization, and overall quality of publications. You can send feedback to DPADDocFeedback@dell.com.



PowerProtect Data Manager File System Agent Overview

Topics:

- PowerProtect Data Manager overview
- Introducing the File System agent
- File System agent system requirements
- Supported Internet Protocol versions
- Encryption in-flight
- PowerProtect Data Manager new deployment overview
- PowerProtect Data Manager existing deployment overview
- Security configuration

PowerProtect Data Manager overview

Use PowerProtect Data Manager with the application agent to perform the following operations:

- Automate the configuration of the application agent backup policy and protection storage settings.
- Create a catalog of backups that the application agent creates. Then, monitor that catalog data to determine if retention policies are being adhered to.
- Manage the life cycle of backups that the application agent creates. Ensure that the backups are marked for garbage collection, based on the rules of the retention policy.

PowerProtect Data Manager does not change the way that the application agent works. DBAs, system administrators, or backup administrators create the backups and perform the restore operations.

Introducing the File System agent

The File System agent enables an application administrator to protect and recover the File System agent application data on the application host. PowerProtect Data Manager integrates with the File System agent to check and monitor backup compliance against protection policies. PowerProtect Data Manager also enables central scheduling for backups.

You can install the File System agent on AIX, Linux, or Windows. See *Enabling the File System Agent*.

NOTE:

PowerProtect Data Manager supports the coexistence of the Microsoft application agent and the File System agent on Windows. When a volume includes any application database and log files:

- File System agent block-based backups of the volume automatically exclude the database and log files from the file system backup.
- File System agent file-based backups of the volume do not automatically exclude the database and log files, but you can explicitly exclude those files through the exclusion filters in the policy. It is recommended that you exclude the application database and log files from file system backups if you use the corresponding application agent to back up the files.

In both cases, File System agent backups do not involve any database writer, regardless of whether or not the database and log files are excluded. The backups do not interfere with the database backup chaining.

To enable the discovery and scheduling of backups with PowerProtect Data Manager, you must approve the client in the PowerProtect Data Manager UI. *Manage the File System agent* provides more information.

Software compatibility information for the PowerProtect Data Manager software and application agents is provided by the E-Lab Navigator.

File System agent system requirements

Ensure that your environment meets the requirements for a new deployment or update of PowerProtect Data Manager.

Requirements:

- NOTE:** The most up-to-date software compatibility information for the PowerProtect Data Manager software and the application agents is provided by the E-Lab Navigator.
- A list of hosts that write backups to DD systems is available.
- DDOS version 6.1 or later and the PowerProtect DD Management Center are required. All models of DD systems are supported.
 - NOTE:** PowerProtect DD Management Center is required with a DDOS version earlier than 6.1.2. With DDOS version 6.1.2 or later, you can add and use a DD system directly without PowerProtect DD Management Center.
- Application agent 19.14 or earlier is required.
- License: A trial license is provided with the PowerProtect Data Manager software. Contact Customer Support for assistance with a permanent PowerProtect Data Manager license.
- Large environments require multiple PowerProtect Data Manager instances. For assistance with sizing requests, contact your Customer Support representative.
- The PowerProtect Data Manager 19.14 download file requires the following:
 - ESXi version 6.5, 6.7, or 7.0.
 - 10 vCPUs, 24 GB RAM, one 100 GB disk, and one 500 GB disk.
 - The latest version of the Google Chrome browser to access the PowerProtect Data Manager UI.
 - TCP port selected as the PowerProtect Data Manager communication port from the supported port ranges 7000 to 7009 and 7012 to 7020. The selected port is open between PowerProtect Data Manager and the application agent host.
 - NOTE:** If a port is not selected from the port ranges, then the default port 7000 is used as the PowerProtect Data Manager communications port.
- VMware ESXi server that hosts PowerProtect Data Manager meets the following minimum system requirements:
 - 10 CPU cores
 - 24 GB of RAM for PowerProtect Data Manager
 - Five disks with the following capacities:
 - Disk 1—100 GB
 - Disk 2—500 GB
 - Disk 3—10 GB
 - Disk 4—10 GB
 - Disk 5—5 GB
 - One 1-GB NIC

Supported Internet Protocol versions

PowerProtect Data Manager and its components support IPv4 and IPv6 addresses in certain configurations.

Table 4. Supported configurations

Component	Internet Protocol
PowerProtect Data Manager core	IPv4 only or both IPv4 and IPv6
PowerProtect Data Manager cloud deployments (AWS, Azure, GCP)	IPv4 only NOTE: Despite other entries in this chart to the contrary, if PowerProtect Data Manager is deployed to a cloud environment, no component in the cloud can use IPv6.
VM Direct, TSDM, and Search	IPv4 only or IPv6 only

Table 4. Supported configurations (continued)

Component	Internet Protocol
	① NOTE: Virtual machines that are backed up must use the same protocol that VM Direct uses. Virtual machines can use both IPv4 and IPv6, even though VM Direct and TSDM cannot.
Application agents integrated with PowerProtect Data Manager:	① NOTE: If both IPv4 and IPv6 are configured and the PowerProtect Data Manager FQDN is used, the agent uses IPv6 for network communication.
• File System	IPv4, IPv6, or both
• Microsoft Exchange Server	IPv4 only or both IPv4 and IPv6
• Microsoft SQL Server (Application Direct)	IPv4, IPv6, or both
• Microsoft SQL Server (VM Direct)	IPv4 only or IPv6 only ① NOTE: Only the Microsoft SQL Server agent supports VM Direct.
• Oracle RMAN	IPv4, IPv6, or both
• SAP HANA	IPv4, IPv6, or both
• Storage Direct	IPv4 only
Stand-alone application agents	IPv4 only
Network-attached storage (NAS)	IPv4, IPv6, or both
Storage arrays (PowerStore)	IPv4 only
Kubernetes	IPv4 only
PowerProtect Data Manager management	IPv4 or IPv6
PowerProtect DD communication	IPv4 or IPv6
Report Browser	IPv4 only ① NOTE: If PowerProtect Data Manager is configured to use both IPv4 and IPv6, configuring an NTP server and setting a time zone is required for accurate date and time information in reports.
SupportAssist	IPv4, IPv6, or both
Syslog Log Server Gateway	IPv4 or IPv6

The following limitations and considerations apply.

Communication with components

If PowerProtect Data Manager is configured to only use one protocol, all components it communicates with must also use that protocol. If some components that PowerProtect Data Manager communicates with use IPv4 and others use IPv6, PowerProtect Data Manager must be configured to use both IPv4 and IPv6.

DD systems and DDVE

If a DD system or a DDVE instance uses only IPv6, the required IPv6 interface must be manually selected when a protection policy is added or edited.

Network-attached storage and DD-system storage units

If the storage unit of a protection policy is different or changed from the destination asset source, you must assign a network to the destination asset for a successful restore. For example, if your source asset is backed up in an IPv6 network, you must assign an IPv6 network to the destination asset for the restore to be successful.

To assign a network for the destination asset, perform the following steps:

1. In the PowerProtect Data Manager UI, select **Infrastructure > Assets > NAS**.
2. Select the destination asset, click **More Actions** and select **Assign Network**. The **Assign Network** page appears.
3. Select a network from the **Network Label** list, click **Save**.
4. If a restore failed because of the wrong destination address, retry the operation.

Disaster recovery

Recovering a PowerProtect Data Manager server might result in a conflict with protection-policy configurations. For instance, if the recovered server is configured to use only IPv4, a protection policy that is configured to use IPv6 cannot run.

Name resolution

Name resolution and reverse IP lookup must be configured to ensure the following:

- Fully qualified domain names of PowerProtect Data Manager, its components, and DD components resolve to a valid IPv4 or IPv6 address.
- If both IPv4 and IPv6 addresses are used for DD, both addresses resolve to the same FQDN.
- All IPv4 and IPv6 addresses are valid and reachable.
- The FQDNs of application-agent hosts that use FQDN as their preferred host address resolve to a valid IPv4 or IPv6 address.
- Each application-agent host that uses FQDN as its preferred host address resolves the FQDN of PowerProtect Data Manager to an IP address of the same protocol that it uses. For example, if a host uses IPv4, it resolves the FQDN of PowerProtect Data Manager to an IPv4 address.

Server updates

IPv6 is only supported with new deployments of PowerProtect Data Manager 19.12 or later. Using IPv6 after updating from PowerProtect Data Manager 19.11 or earlier is unsupported.

Search Engine indexing and adding IPv6 to an IPv4-only system

If you add IPv6 to an IPv4-only system, indexing from any existing Search Engine cluster becomes unavailable. After adding IPv6, you must delete all IPv4 Search Engine nodes to remove the Search Engine cluster, and then add new IPv6 nodes to a new cluster.

Unlike other PowerProtect Data Manager components, if IPv6 is used with a Search Engine, the FQDN of all Search Engine nodes and related DD systems must always resolve to an IPv6 address and never to an IPv4 address.

Storage Policy Based Management

If using vCenter or ESXi 7.0u2 or earlier with only IPv6, SPBM providers must be added using their PowerProtect Data Manager FQDN.

Service Unavailable messages with the vSphere Client PowerProtect plug-in

If vCenter uses the vSphere Client PowerProtect plug-in with IPv6 and the vCenter host is added to PowerProtect Data Manager using its IPv6 address or FQDN, `Service Unavailable` messages might be seen for the protected virtual machine. Backups and restores of the protected virtual machine are unaffected, and these messages can be ignored.

Uncompressed IPv6 formatting

Network interfaces that exist on a DD 7.4.x or earlier system and that are configured to use an uncompressed IPv6 format cannot be discovered. An example of an uncompressed IPv6 format is `2620:0000:0170:0597:0000:0000:0001:001a`. An example of a compressed IPv6 format is `2620:0:170:597::1:1a`. To use these network interfaces, reconfigure them to use either an IPv4 address or a compressed IPv6 address, and then initiate a discovery.

Encryption in-flight

PowerProtect Data Manager provides centralized management of encryption in-flight for supported workloads. Encryption in-flight is supported for both centralized and self-service policies, where applicable.

You can ensure that backup and restore content is encrypted when read on the source, transmitted in encrypted form, and then decrypted before it is saved on the destination. This prevents another party from intercepting private data.

PowerProtect Data Manager only supports encryption in-flight for File System, Kubernetes clusters, Microsoft SQL Server, Microsoft Exchange Server, network-attached storage (NAS), PowerStore storage arrays, Oracle, SAP HANA, and VMware virtual machine workloads. This is a global setting that is applicable to all supported workloads.

For File System, Microsoft Exchange Server, Oracle, SAP HANA, and NAS workloads, encryption in-flight is only supported for Application Direct hosts. For File System agents, restore encryption is supported for image-level restore only. For Microsoft SQL Server agents, restore encryption is supported for database-level restore only.

The *PowerProtect Data Manager Administrator Guide* and *PowerProtect Data Manager Security Configuration Guide* provide more information about encryption in-flight, such as how to enable the feature and important considerations to understand before enabling.

PowerProtect Data Manager new deployment overview

Familiarize yourself with the high-level steps required to deploy PowerProtect Data Manager with the File System agent.

Steps

1. Design how to group the backups based on the storage requirements and retention policies.
The account team can help with backup storage design.
2. Install PowerProtect DD Management Center (DDMC).
PowerProtect Data Manager uses DDMC to connect to the DD systems. The *DD Management Center Installation and Administration Guide* provides instructions.
3. Deploy PowerProtect Data Manager from the download file.
The *PowerProtect Data Manager Deployment Guide* provides instructions.
4. Add external DD systems or DDMC to PowerProtect Data Manager.
The *PowerProtect Data Manager Administration and User Guide* provides instructions on how to add protection storage.
5. Install the File System agent on the appropriate hosts and connect them to PowerProtect Data Manager according to the instructions in the next chapter.
6. Add new or approve pending agent requests in the PowerProtect Data Manager according to the instructions in the next chapter.
7. Add a protection policy for groups of assets that you want to back up.

NOTE: After you create a centralized protection job, the first backup is a full backup.

The *PowerProtect Data Manager Administration and User Guide* provides instructions.

8. Add Service Level Objectives to the protection policy to verify that the protected assets meet the service-level agreements (SLAs).

The *PowerProtect Data Manager Administration and User Guide* provides instructions.

9. Now that the configuration is complete, it is recommended to perform a full backup so that PowerProtect Data Manager can detect the proper backup chain.

Without a full backup, PowerProtect Data Manager treats the backups as partial and assumes that you are out of compliance.

10. Monitor protection compliance in the PowerProtect Data Manager dashboard.

PowerProtect Data Manager existing deployment overview

Familiarize yourself with the high-level steps required to deploy PowerProtect Data Manager with the File System agent to an existing environment.

Steps

1. Install PowerProtect DD Management Center (DDMC).

PowerProtect Data Manager uses DDMC to connect to the DD systems. The *DD Management Center Installation and Administration Guide* provides instructions.

2. Deploy PowerProtect Data Manager from the download file.

The *PowerProtect Data Manager Deployment Guide* provides instructions.

3. Add external DD systems or DDMC to PowerProtect Data Manager.

The *PowerProtect Data Manager Administration and User Guide* provides instructions on how to add protection storage.

4. Update the File System agent, or uninstall and then reinstall the agent on the hosts, and connect them to PowerProtect Data Manager according to the instructions in the next chapter.

5. Add new or approve pending agent requests in the PowerProtect Data Manager according to the instructions in the next chapter.

6. Add a protection policy for groups of assets that you want to back up.

NOTE: After you create a centralized protection job, the first backup is a full backup.

The *PowerProtect Data Manager Administration and User Guide* provides instructions.

7. Add Service Level Objectives to the protection policy to verify that the protected assets meet the Service Level Agreements (SLAs).

The *PowerProtect Data Manager Administration and User Guide* provides instructions.

8. Now that the configuration is complete, it is recommended to perform a full backup so that PowerProtect Data Manager can detect the proper backup chain.

Without a full backup, PowerProtect Data Manager treats the backups as partial and assumes that you are out of compliance.

9. Monitor protection compliance in the PowerProtect Data Manager dashboard.

Security configuration

A separate guide provides some server configuration tasks which are intended specifically for PowerProtect Data Manager security administrators, whose role may be separate from the PowerProtect Data Manager host system administrator.

The *PowerProtect Data Manager Security Configuration Guide* provides detailed instructions for all security-related tasks, including but not limited to:

- Port requirements for and between the following components:

- PowerProtect Data Manager
- Configured DD systems
- VM Direct appliances (embedded and external)
- Application-agent hosts
- Web and REST API clients
- Callhome (SupportAssist)
- ESXi
- vCenter
- Configuring identity providers
- Managing local and external user accounts
- Changing and resetting passwords
- Assigning users and groups to roles and associated privileges
- Managing credentials for local and remote components
- Creating resource groups to define scopes of authority
- Managing security certificates, where applicable

Role-based security

PowerProtect Data Manager provides predefined user roles that control access to areas of the user interface and to protected operations. Some PowerProtect Data Manager functionality is reserved for particular roles and may not be accessible from every user account.

By using the predefined roles, you can limit access to PowerProtect Data Manager and to backup data by applying the principle of least privilege.

The *PowerProtect Data Manager Security Configuration Guide* provides more information about user roles, including the associated privileges and the tasks that each role can perform.

Enabling the File System Agent

Topics:

- About the File System agent
- Application agent and File System agent coexistence
- File System agent prerequisites
- File System agent limitations
- Block-based backups
- Best practices for file system backups
- Configure the file system parallel backup setting
- Configure asset multi-streaming for file-based backups
- Protect an asset with the File System agent
- Protect an asset in a Microsoft Windows Server clustered environment with the File System agent
- Installing and uninstalling the File System agent on AIX
- Installing and updating the File System agent on Linux
- Installing and updating the File System agent on Windows
- Manage the File System agent
- Configure the centralized file-level restore service port number
- Update the application agent in the PowerProtect Data Manager UI
- Moving a File System agent from one PowerProtect Data Manager system to another
- Enable the File System agent after migrating File System assets to a refreshed server

About the File System agent

The File System agent enables an application administrator to protect and recover data on the file system host. PowerProtect Data Manager integrates with the File System agent to check and monitor backup compliance against protection policies. PowerProtect Data Manager also enables central scheduling for backups.

You can install the File System agent on the host that you plan to protect by using the installation wizard. Installing and uninstalling the File System agent on AIX, installing and updating the File System agent on Linux, and installing and updating the File System agent on Windows provide instructions.

Software compatibility information for the PowerProtect Data Manager software and the File System agent is provided in the E-Lab Navigator.

Application agent and File System agent coexistence

PowerProtect Data Manager supports the following application agent and File System agent coexistence:

- Coexistence of the Oracle RMAN agent or SAP HANA agent with the File System agent on Linux.
- Coexistence of the Oracle RMAN agent with the File System agent on AIX.
- Coexistence of the Microsoft SQL Server or Microsoft Exchange Server application agent with the File System agent on Windows.

NOTE: When the Microsoft Exchange Server application agent and the File System agent coexist and both agents are installed and registered to the same PowerProtect Data Manager instance, the following workflows are supported.

For File Systems:

- Use File-based backup (FBB) instead of Block-based backup, and provide a dummy exclusion filter in the protection policy.
- In the File System protection policy backup, do not include Microsoft Exchange Server.edb and log file assets.

For Microsoft Exchange Server:

- o Run Microsoft Exchange Server full backups only.

The coexistence of these agents enables you to protect the Microsoft SQL Server, Microsoft Exchange Server, Oracle, or SAP HANA database with the host file system. The following configurations are supported for agent coexistence:

- Both agents in managed mode (registered to PowerProtect Data Manager)
- The Microsoft SQL Server, Microsoft Exchange Server, Oracle, or SAP HANA agent in stand-alone mode, with the File System agent registered to PowerProtect Data Manager

NOTE: The latest version of each agent must be installed if the agents are registered to PowerProtect Data Manager. The File System agent is supported in managed mode only.

The steps for installation and usage for each agent are the same.

The table below lists the supported use cases and limitations.

Table 5. Supported cases

Category	Supported cases	Current limitations
Agent installation and uninstallation	<ul style="list-style-type: none"> • New installation of both agents is supported with: <ul style="list-style-type: none"> o Microsoft SQL Server agent, Microsoft Exchange Server agent, Oracle RMAN agent, or SAP HANA agent in stand-alone or managed mode. o File System agent in managed mode. • New installation of an agent is supported in managed mode with an existing agent in stand-alone mode. • New installation of an agent is supported in stand-alone mode with an existing agent in managed mode. • Repair of an existing agent installation is supported. • Uninstallation of agents is supported. 	<ul style="list-style-type: none"> • Uninstalling the last agent that is installed on the host unregisters the host from PowerProtect Data Manager. Any new agent installation that occurs after the uninstall must be newly registered to the PowerProtect Data Manager server. • Similar to the agent installations, uninstallation of each agent is performed separately.
Host registration and unregistration	<ul style="list-style-type: none"> • Registration of an installed agent to the PowerProtect Data Manager server is supported. • Unregistration of agents from the PowerProtect Data Manager server is supported. 	<ul style="list-style-type: none"> • Both agents, if operating in managed mode, should be registered to the same PowerProtect Data Manager server only. There is no option to register each agent to a different PowerProtect Data Manager server. • Unregistering a host unregisters all the managed agents that are installed on that host. Stand-alone agents are not affected. • After unregistering a host, the host's assets still appear in the UI in order to support the restore of these assets to a different host. However, backups are not initiated on these assets as the protection policies are disabled.
Backup and restore features	<ul style="list-style-type: none"> • Protection policy creation is supported on all registered agents. • All scheduled protection policy backups are supported on both agents as per individual protection policies. • Self-service backups are supported on both agents. • Compliance is supported on both agents as per the individual service-level agreements (SLAs). • Manual backups are supported at the protection policy level and individual asset level through the centralized protection policy workflow. 	

File System agent prerequisites

Review the following prerequisites before installing and using the File System agent in PowerProtect Data Manager.

Windows, Linux, and AIX prerequisites

- Ensure that your host is a 64-bit system. PowerProtect Data Manager supports only 64-bit hosts.
- Ensure that your host is a supported operating system version. Software compatibility information for the PowerProtect Data Manager software is provided in the E-Lab Navigator.
- Ensure that all clocks on both the host and PowerProtect Data Manager are time-synced to the local NTP server to ensure discovery of the backups.
- Ensure that the host and PowerProtect Data Manager can see and resolve each other.

If PowerProtect Data Manager and the File System agent are registered in different domains, add IP address and Fully Qualified Domain Name (FQDN) entries to the hosts files on both the client and server:

IP address	FQDN	common name
10.10.100.100	yourdomain.com	yourdomain

- Windows: C:\Windows\System32\drivers\etc\hosts
- Linux: /etc/hosts
- LVM and VxVM partitions or volumes are supported.
- Physical or non-LVM partitions are supported only when using file-based backups.
- Each volume group on LVM2 or VxVM must have at least 10% free space for a block-based backup to succeed. For successful Windows VSS snapshot during block-based or file-based backup, the free space requirement is 20%.
- For file-based backups, ensure that the drive on which the File System agent is installed has adequate free space for the metadata record files that are created during a backup. Provide about 250 MB free space for each million files that you are backing up.
- For any ESXi version 6.5 or earlier host with PowerStore storage attached, the Windows operating system deployment or installation cannot proceed. If the DiskMaxIOSize parameter is not configured with the proper value, the File System agent backup and restore operations fail. Ensure that you set the DiskMaxIOSize to 1024 KB.
- Ensure that no service is using port 7010 or 7011. These ports must be reserved for File System agent operations.
- Review the limitations in the section File System agent limitations.

Additional Linux prerequisites

- File system discovery requires an ext3, ext4, XFS, or BTRFS file system type.
- If there is no block-based backup (BBB) driver available on the Linux host due to kernel version mismatch, you can still perform file-based backups. Install the File System agent by using the `install.sh --skip-driver` command instead.
- On the Linux hosts that have the UEFI Secure Boot option enabled, block-based backup drivers do not load, and the error message `insmod: ERROR: could not insert module /lib/modules/3.10.0-693.el7.x86_64/extra/nrbbb.ko: Required key not available` appears. As a workaround, you can disable the Secure Boot option.
- Ensure that the file system has an entry in `/etc/fstab`. Otherwise, discovery fails.
- To enable file-level restores, complete the following:
 1. Log in to the system you are restoring from as `root`.
 2. Install iSCSI client packages.

See the Operating System documentation for the installation procedure.
 3. For **Service Start**, choose **Manually**, and then click **OK**.
- If installing the block-based backup driver, review the output of `/proc/sys/kernel/kptr_restrict` to verify that permissions to install the driver are set. If the value is set to `2`, then a restriction exists that might result in block-based backup driver installation failure. Run the following command to change this setting: `echo 1 > /proc/sys/kernel/kptr_restrict`
- Install the `lsb_release` package. The operating system documentation provides package installation procedures.

Additional AIX prerequisites

- File system discovery requires a JFS or JFS2 file system type.
- Ensure that IBM XL C or C++ Runtime for AIX 16.1.0.7 and later is installed.

File System agent limitations

Review the following limitations that are related to File System agent support in PowerProtect Data Manager.

Software compatibility

Software compatibility information for the PowerProtect Data Manager software and the File System agent is provided in the E-Lab Navigator.

Windows, Linux, and AIX limitations

- File-based backup (FBB) cannot be performed in parallel for a large file system when the entire data is present in a single folder. Because, instead of creating multiple chunks, the system creates one chunk causing the backup runs through a single stream.

Windows and Linux limitations

- File System agent block-based backups exclude the following:
 - Application files such as Microsoft SQL Server and Microsoft Exchange Server files.
 - i NOTE:** For file-based backups, application data such as Microsoft SQL Server and Microsoft Exchange Server data is backed up, but you can explicitly exclude the data through the exclusion filters in the policy. It is recommended that you exclude the application data from file system backups if you use the corresponding application agent to back up the data. File System agent backups do not involve any database writer, regardless of whether or not the application data is excluded. The backups do not interfere with the database backup chaining.
 - HyperVisor files. You cannot run a Hyper-V server in a VMware virtual machine.
 - Data belonging to individual application writers.
 - Unsupported application writer files.
- It is recommended to use different mount points for each drive. Reusing mount points might cause unexpected issues during file system discovery.
- If a Windows or Linux file system host is unregistered from PowerProtect Data Manager and then reregistered with a different FQDN, because PowerProtect Data Manager recognizes the registration as a new host by its new name, duplicate asset entries appear in the UI—those for the host that is registered earlier, as well as for the host that is registered by the new name. This does not impact backup and restore functionality on the new host.
- The File System agent and application agents use the FQDN for registration. If the File System agent coexists with the Microsoft SQL, Oracle, or SAP HANA application agent, both agents must use the FQDN.
- For a protection policy backup with assets from different hosts, the backup status appears as *Failed* in the UI if the backup of one asset within the policy fails.
- Running the `ddfsrv` and `ddfsrc` commands to perform self-service backup and restore of file systems fails if you provide the DD hostname (instead of IP) for the `DFA_Si_DD_HOST` variable.
- PowerProtect Data Manager does not support block-based image recoveries if the sector sizes (in bytes) for the source and target volumes are different. For example, you cannot perform a block-based image recovery of a volume with 4096-byte sectors to a volume with 512-byte sectors, and vice versa.

Windows Limitations

- The file-level restore of a folder can result in the loss of the sparse flag of any sparse files within the folder. To preserve the sparse flag of these files, restore the files individually.

Configure asset multi-streaming for file-based backups

PowerProtect Data Manager supports asset multi-streaming, which enables you to run a file system backup of an asset in parallel streams to reduce the time required to complete the file-based backup. Asset multi-streaming is enabled by default.

When using asset multi-streaming for file-based backups of large volumes, the contents of the volumes are divided into chunks. These chunks are backed up in parallel in multiple streams to increase the backup throughput.

The chunk size is a configurable parameter, with a default value of 50 GB (`--chunksize=50`). The number of parallel streams is set by the `--max-host-streams` parameter in the configuration file, or the `-M` command-line option, as described in *Configure the file system parallel backup setting*. The value for number of parallel streams set in the config file takes precedence over the value set in the CLI. The number of streams is set per host (and not per asset), and each asset uses the streams available to it.

Creating chunks and multi-streaming the backup consumes some extra computing resources. You can control that usage by tuning the `chunksize` and `--max-host-streams` parameters. If, after tuning, you are still not satisfied with the resource usage, you can disable multi-streaming. To do so, add the `--disable-asset-multistreaming=true` parameter in the configuration file, `.ddfsrv.fsagentconfig` under the `Settings` folder.

Multi-streaming may not yield desired results in the following scenarios:

- System has limited CPUs or memory.
- Reads are slow, in which case, CPU usage will be high. For disks with slow read or high latency, disable multi-streaming.
- You are backing up many small files. Multi-streaming gives better performance when files are large. Many small files can cause undesirably high CPU usage levels.

Recommendations for optimal performance:

- Multi-streaming can be CPU-intensive. It is important to tune the environment for optimal chunk size and an optimal number of streams.
- If there are more volumes than streams in a protection policy, disable asset multi-streaming.

Protect an asset with the File System agent

The following task describes the steps required to protect an asset with a protection policy.

Steps

1. Add a storage system.
For more information, see the *PowerProtect Data Manager Administrator Guide*.
2. Install the File System agent on the file system host.
For more information, see the section on installing the File System agent on the operating system of the host.
3. Add or approve the File System agent on the file system host.
For more information, see *Manage the File System agent*.
4. Discover the file system asset.
For more information, see *Discover a file system host*.
5. Create a protection policy to protect the file system.
For more information, see *Protection policies for File System agent*.

NOTE: You cannot perform a backup to a secondary DD system. You can only restore from a secondary DD system.

Protect an asset in a Microsoft Windows Server clustered environment with the File System agent

The following task describes the steps required to protect clustered disks and Cluster Shared Volumes with a protection policy.

About this task

Repeat these steps for each node in the cluster that is registered with PowerProtect Data Manager.


Steps

1. Add a storage system.

For more information, see the *PowerProtect Data Manager Administrator Guide*.

2. Install the File System agent on the node.

For more information, see *Install the File System agent on Windows*.


 **NOTE:** Cluster assets are automatically discovered after the agent is installed.

3. Add or approve the File System agent on the node.

For more information, see *Manage the File System agent*.


4. Discover the file system asset.

For more information, see *Discover a file system host*.


 **NOTE:** Stand-alone assets on a node are listed under the name of the cluster node. Cluster assets on a node are listed under the name of the logical cluster host.

5. Create a protection policy to protect the cluster.

For more information, see *Protection policies for File System agent*.

 **NOTE:** The backup of a cluster asset is routed through the node on which the asset or volume is active.

6. If you are planning to use bare-metal recovery for disaster recovery, record the hardware configuration of the cluster-node Windows clients.

 **CAUTION:** Bare-metal recovery can fail if the hardware configuration of the new Windows clients does not match the hardware configuration of the old cluster-node Windows clients.

Necessary information includes the hardware vendor, size and type of disks, type of network adapters, and amount of memory assigned.

Installing and uninstalling the File System agent on AIX

Learn how to install and uninstall the File System agent on AIX.

Install the File System agent on AIX

Install the File System agent on supported AIX systems using an interactive or silent installation procedure.

Install the File System agent on AIX in interactive mode

Use this interactive procedure to install the File System agent on supported AIX systems.

Prerequisites


- Ensure that you review the prerequisites provided in *File System agent prerequisites*.
- Ensure that the AIX host has sufficient free space to install the File System agent.

- Ensure that there is no endpoint security or monitoring software such as Carbon Black installed on the AIX host, which might prevent the File System agent installation.
- Download the File System agent software package to the AIX host.

NOTE: If a value is not provided for *PowerProtect Server IP* in the installation commands, the product is installed without PowerProtect registration, and no backups can be initiated from the UI.

Steps

1. In the PowerProtect Data Manager UI:

- Click  and then select **Downloads** from the **System Settings** menu.
- Select the File System agent download package for AIX, `fsagent19x_aixpower72.tar.gz`.
- Download the package in the location that you want to install the File System agent.

NOTE: Relocating the installation to another partition or mount point on AIX is not supported.

2. Untar the installer by running the following commands:

- `gunzip fsagent19x_aixpower72.tar.gz`
- `tar -xvf fsagent19x_aixpower72.tar`

3. To change the current working directory to the extracted path, run the command `cd fsagent`.

4. Run the installation script `install.sh`.

To run in debug mode, run `install.sh --debug`.

To get help, run `install.sh --help`.

NOTE: File System agent does not support block-based backups on AIX. All backups performed by File System agent on AIX will be file-based backups.

The following `.rte` files are installed as part of the script:

- `powerprotect-agent.svc.rte` —Installs or updates the agent service component for the File System agent.
- `ppdm_fsagent.rte` —Installs the File System agent related files and folders.

5. Type the PowerProtect Data Manager server FQDN or IP address. It is recommended to use the FQDN.

NOTE:

- If another application agent is already installed on the client and registered to PowerProtect, ensure that you register the agent with the existing PowerProtect Data Manager server FQDN. When you register the agent with a PowerProtect Data Manager server that is different from the currently registered server, no warning message appears, and requests are routed to the newer server instance.
- If you specify a hostname or fully qualified domain name (FQDN) with an underscore (`_`) for the PowerProtect Data Manager server, then the communication will be done by the system's IP, if provided to the system on registration.

6. Type the preferred FQDN or IP address of the application host.

7. Type the port number from the supported port ranges 7000 to 7009 and 7012 to 7020. The specified port is used for communication between the File System agent and PowerProtect Data Manager.

NOTE: If you do not specify a port number, the default port 7000 is used as the communication port. The ports 7010 and 7011 are used by the agent service message bus.

8. Type the answer to the prompt about whether to add the firewall rule for the specified port. If you answer **yes**, the `/opt/dpsapps/agent.svc/configfw.sh` script runs to set the firewall rule and enable the PowerProtect Data Manager communication port.

9. To enable the PowerProtect Data Manager communication TCP port manually, run the `/opt/dpsapps/agent.svc/configfw.sh` script as the root user.

The system responds that the firewall script is running, and is configuring firewall rules.

Install the File System agent on AIX in silent mode

On AIX, review the following commands to perform a silent installation of the File System agent.

- ① **NOTE:** The `--server` option is used to specify the PowerProtect Data Manager server IP for registration, and is mandatory for silent install.

Use the following commands to perform a silent installation on AIX:

- For a silent installation, including registration of the agent with the PowerProtect Data Manager server, type: `install.sh --silent-install --server=<PowerProtect_Data_Manager_server_IP>`
- To run the installer in debug mode during silent installation, type: `install.sh --debug --silent-install --server=<PowerProtect_Data_Manager_server_IP>`
- To skip installation of the block-based backup driver, type: `install.sh --skip-driver --silent-install --server=<PowerProtect_Data_Manager_server_IP>`
- To specify the FQDN, IPv4 address, or IPv6 address of the File system agent host to be registered, and the port to be used for communication between the agent and PowerProtect Data Manager, type: `install.sh --silent-install --server=<PowerProtect_Data_Manager_server_IP> --preferred-address=<agent_host_FQDN_or_IP> --client-port=<port_number>`

- ① **NOTE:** The port number must be from the supported port ranges 7000 to 7009 and 7012 to 7020. If you do not specify a port number, the default port 7000 is used as the communication port. The ports 7010 and 7011 are used by the agent service message bus.

- To enable the PowerProtect Data Manager communication TCP port, run the `/opt/dpsapps/agentsvc/configfw.sh` script as the root user.

Uninstall the File System agent on AIX

On AIX, you can uninstall the File System agent by performing the following steps:

Steps

1. Obtain the `uninstall.sh` script from the folder extracted from the package `fsagent19x_aixpower72.tar.gz` or under `/opt/dpsapps/fsagent/bin`.
2. Run `./uninstall.sh`.

The following message appears:

```
Other application agents might be using powerprotect-agentsvc. Do you wish to
uninstall powerprotect-agentsvc? [y/n]
```

3. To confirm that you want to uninstall `powerprotect-agentsvc`, type **Y**.

- ① **NOTE:** If you type **N**, the `.rte` files (`ppdm_fsagent.rte`, `ppdm-bbbwt.rte` and `ppdm-fsagent.rte`) are uninstalled. However, `powerprotect-agentsvc` remains in an installed state.

There is no silent uninstall procedure on AIX.

Installing and updating the File System agent on Linux

Learn how to install and update the File System agent on Linux.

Install the File System agent on Linux

Install the File System agent on supported Linux systems using an interactive or silent installation procedure.

Install the File System agent on Linux in interactive mode


Use this interactive procedure to install the File System agent on supported Linux systems.

Prerequisites

- Ensure that you review the prerequisites provided in File System agent prerequisites.
 - Download the File System agent software package to the Linux host.
- ① **NOTE:** If a value is not provided for *PowerProtect Server IP* in the installation commands, the product is installed without PowerProtect registration, and no backups can be initiated from the UI.

Steps

1. In the PowerProtect Data Manager UI:

- a. Click , and then select **Downloads** from the **System Settings** menu.
- b. Select the File System agent download package for Linux, `fsagent19x_linux_x86_64.tar.gz`.
- c. Download the package in the location that you want to install the File System agent.

① **NOTE:** Relocating the installation to another partition or mount point on Linux is not supported.

2. Untar the installer by running `tar -xvf fsagent19x_linux_x86_64.tar.gz`.

Then, run the command `cd fsagent` to change the current working directory to the extracted path.

① **NOTE:** To verify the authenticity and integrity of the RPM files prior to the installation step, follow the instructions in the *PowerProtect Data Manager Security Configuration Guide*.

3. Run the installation script `install.sh`.

To run in debug mode, run `install.sh --debug`.

To get help, run `install.sh --help`.

① **NOTE:** For installations on Oracle Linux distributions or CentOS Linux distributions (for CentOS 8.0, 8.1, and 8.2), run `install.sh --skip-driver` to skip the block-based backup driver installation. These distributions do not currently support block-based backups. All backups performed by the File System agent on these distributions will be file-based backups.

For RHEL distributions (for Red Hat 8.0, 8.1, and 8.2), Security-Enhanced Linux (SELinux) is enabled by default. It can support block-based backups, provided you continue the installation with the procedure in *Install the File System agent on RHEL distributions*.

The following `.rpm` or `.deb` files are installed as part of the script:

- `powerprotect-agentsvc.rpm` or `powerprotect-agentsvc.deb`—Installs or updates the agent service component for the File System agent.
- `ppdm_bbbwt.rpm` or `ppdm-bbbwt.deb`—Installs the block-based backups driver.
- `ppdm_fsagent.rpm` or `ppdm-fsagent.deb`—Installs the File System agent related files and folders.

4. Type the PowerProtect Data Manager server FQDN or IP address. It is recommended to use the FQDN.

① **NOTE:**

- If another application agent is already installed on the client and registered to PowerProtect, ensure that you register the agent with the existing PowerProtect Data Manager server FQDN. When you register the agent with

a PowerProtect Data Manager server that is different from the currently registered server, no warning message appears, and requests are routed to the newer server instance.

- If you specify a hostname or fully qualified domain name (FQDN) with an underscore (_) for the PowerProtect Data Manager server, then the communication will be done by the system's IP, if provided the system on registration.

5. Type the preferred FQDN or IP address of the application host.

6. Type the port number from the supported port ranges 7000 to 7009 and 7012 to 7020. The specified port is used for communication between the File System agent and PowerProtect Data Manager.

NOTE: If you do not specify a port number, the default port 7000 is used as the communication port. The ports 7010 and 7011 are used by the agent service message bus.

7. Type the answer to the prompt about whether to add the firewall rule for the specified port. If you answer **yes**, the `/opt/dpsapps/agentavc/configfw.sh` script runs to set the firewall rule and enable the PowerProtect Data Manager communication port.

8. To enable the PowerProtect Data Manager communication TCP port manually, run the `/opt/dpsapps/agentavc/configfw.sh` script as the root user.

The system responds that the firewall script is running, and is configuring firewall rules.

NOTE: If the firewall rules are not applied, restart the firewall.

Next steps

If the host is not already approved, add the file system host to the PowerProtect Data Manager server. Manage the File System agent provides more information.

Discover file system assets. Discover a file system host provides more information.

Install the File System agent on RHEL distributions

For RHEL distributions (for Red Hat 8.0, 8.1, and 8.2), Security-Enhanced Linux (SELinux) is enabled by default. To support block-based backups, run this special installation procedure.

Prerequisites

Complete steps 1 and 2 in Install the File System agent on Linux in interactive mode.

Steps

1. If you have not yet run the installation script, `install.sh`, run it now.

Installation of the block-based backup driver fails.

2. Check that an error message similar to the following appears in `/var/log/messages`

```
insmod: ERROR: could not insert module /lib/modules/4.18.0-80.el8.x86_64/extra/nrbbb.ko:
Permission denied.
```

3. To check the audit log, run: `ausearch -o 'insmod'`

```
It returns a string similar to: type=AVC msg=audit(1624349147.478:628):
avc: denied { module_load } for pid=80964
comm='insmod' path="/opt/dpsapps/fsagent/bin/nrbbb-redhatenterprise-8.2-4.18.0-193.ko"
dev="dm-0" ino=12098527 scontext=system_u:system_r:unconfined_service_t:s0
tcontext=unconfined_u:object_r:bin_t:s0 tclass=system permissive=0
```

`type=AVC` indicates that the installation of the block-based backup driver is failing due to the SELinux policy.

4. To change the SELinux policy so that it will be able to allow access to the block-based backup driver, run: `ausearch -o 'insmod' --raw | audit2allow -M ppdm-fsagent`

It generates two files in the current directory: `ppdm-fsagent.pp` and `ppdm-fsagent.te`

5. To apply the SELinux policy changes, to enable access to the block-based backup driver, run: `semodule -i ppdm-fsagent.pp`

6. Run the installation script once again: `install.sh`

Installation of the block-based backup driver should succeed, and the following `.rpm` or `.deb` files are installed as part of the script:

- `powerprotect-agentsvc.rpm` or `powerprotect-agentsvc.deb`—installs or updates the agent service component for the File System agent.

- `ppdm_bbbwt.rpm` or `ppdm-bbbwt.deb`—Installs the block-based backups driver.
 - `ppdm_fsagent.rpm` or `ppdm-fsagent.deb`—Installs the File System agent related files and folders.
7. Type the PowerProtect Data Manager server FQDN or IP address. It is recommended to use the FQDN.

NOTE: If another application agent is already installed on the client and registered to PowerProtect, ensure that you register the agent with the existing PowerProtect Data Manager server FQDN. When you register the agent with a PowerProtect Data Manager server that is different from the currently registered server, no warning message appears, and requests are routed to the newer server instance.
 8. Type the preferred FQDN or IP address of the application host.
 9. Type the port number from the supported port ranges 7000 to 7009 and 7012 to 7020. The specified port is used for communication between the File System agent and PowerProtect Data Manager.

NOTE: If you do not specify a port number, the default port 7000 is used as the communication port. The ports 7010 and 7011 are used by the agent service message bus.
 10. Type the answer to the prompt about whether to add the firewall rule for the specified port. If you answer **yes**, the `/opt/dpsapps/agentsvc/configfw.sh` script runs to set the firewall rule and enable the PowerProtect Data Manager communication port.

Next steps

If the host is not already approved, add the file system host to the PowerProtect Data Manager server. Manage the File System agent provides more information.

Discover file system assets. Discover a file system host provides more information.

Install the File System agent on Linux in silent mode

On Linux, review the following commands to perform a silent installation of the File System agent.

NOTE: The `--server` option is used to specify the PowerProtect Data Manager server IP for registration, and is mandatory for silent install.

To verify the authenticity and integrity of the RPM files prior to the installation step, follow the instructions in the *PowerProtect Data Manager Security Configuration Guide*.

Use the following commands to perform a silent installation on Linux:

- For silent installation, including registration of the agent with the PowerProtect Data Manager server, type: `install.sh --silent-install --server=<PowerProtect_Data_Manager_server_IP>`
- To run the installer in debug mode during silent installation, type: `install.sh --debug --silent-install --server=<PowerProtect_Data_Manager_server_IP>`
- To skip installation of the block-based backup driver, type: `install.sh --skip-driver --silent-install --server=<PowerProtect_Data_Manager_server_IP>`
- To specify the FQDN, IPv4 address, or IPv6 address of the File system agent host to be registered, and the port to be used for communication between the agent and PowerProtect Data Manager, type: `install.sh --silent-install --server=<PowerProtect_Data_Manager_server_IP> --preferred-address=<agent_host_FQDN_or_IP> --client-port=<port_number>`

NOTE: The port number must be from the supported port ranges 7000 to 7009 and 7012 to 7020. If you do not specify a port number, the default port 7000 is used as the communication port. The ports 7010 and 7011 are used by the agent service message bus.

- To enable the PowerProtect Data Manager communication TCP port, run the `/opt/dpsapps/agentsvc/configfw.sh` script as the root user.

Update the File System agent on Linux

The File System agent supports a direct update from an earlier version if you are using an earlier version of PowerProtect Data Manager. You can update the PowerProtect Data Manager File System agent on supported Linux systems in silent mode.

Update and register the latest version of the PowerProtect Data Manager File System agent for Linux with the same PowerProtect Data Manager server in the same location.

① **NOTE:** A reboot is not required after the completion of the software update.

To update the File System agent on Linux, run the following command:

```
install.sh
```

- For silent update in debugging mode, type: **install.sh --force-upgrade --debug**

① **NOTE:** During silent installation or update, the debug output is directed to syslog to be viewed later.

- To enable the PowerProtect Data Manager communication 7000 TCP port, run the `/opt/dpsapps/agentsvc/configfw.sh` script as the root user.

① **NOTE:**

- When you install or update the File System agent, other app agents on the system must be updated to the same version as the File System agent.
- Following an update, the first block-based backup is promoted to a full backup.

Uninstall the File System agent on Linux

On Linux, you can uninstall the File System agent by performing the following steps:

Steps

1. Obtain the `uninstall.sh` script from the folder extracted from the package `fsagent19x_linux_x86_64.tar.gz` or under `/opt/dpsapps/fsagent/bin`.
2. Run `./uninstall.sh`.

The following message appears:

```
Other application agents might be using powerprotect-agentsvc. Do you wish to
uninstall powerprotect-agentsvc? [y/n]
```

3. Type **y** to confirm that you want to uninstall `powerprotect-agentsvc`.

① **NOTE:** If you type **n**, the `.rpm` files (`ppdm_bbbwt.rpm` and `ppdm_fsagent.rpm`) and `.deb` files (`ppdm-bbbwt.deb` and `ppdm-fsagent.deb`) are uninstalled. However, `powerprotect-agent` remains in an installed state.

There is no silent uninstall procedure on Linux.

Reregister the File System agent on Linux or AIX

You can use the procedure in this topic to reregister the deleted File System agent to the same PowerProtect Data Manager server.

① **NOTE:** You can run the `install.sh` or `register.sh` script to register and reregister the File System agent with PowerProtect Data Manager only if you have not uninstalled the File System agent and PowerProtect agent service.

If you have cleaned up the installation directories and manually uninstalled both the File System agent and PowerProtect agent service, then you must complete the installation procedures in Installing and updating the File System agent on Linux.

To use the `install.sh` script to register and reregister the File System agent with PowerProtect Data Manager, run:

```
install.sh --server=10.125.19.40 --debug
```

The message `AgentService is recommissioned` is displayed to confirm that the agent has been successfully reregistered.

To use the `register.sh` script to register and reregister the File System agent with PowerProtect Data Manager, run:

```
/opt/dpsapps/agentsvc/register.sh --enable
```

Installing and updating the File System agent on Windows

Learn how to install and update the File System agent on Windows.

Install the File System agent on Windows

Install the File System agent on supported Windows systems using an interactive or silent installation procedure.

- NOTE:** In a clustered environment, install the same version of the File System agent on each node in the cluster that is registered with PowerProtect Data Manager.

Install the File System agent on Windows in interactive mode

Use this interactive procedure to install the File System agent on supported Windows systems.

Prerequisites

- Ensure that you carry out the prerequisites provided in File System agent prerequisites.
- Download the File System agent software package.

Steps

1. In the PowerProtect Data Manager UI:
 - a. Click , and then select **Downloads** from the **System Settings** menu.
 - b. Select the File System agent download package for Windows, `fsagent19x_win_x64.zip`.
 - c. Download the package in the location that you want to install the File System agent.
2. Run the `fsagent-19.x.x.x.exe` program.
3. Follow the wizard installation steps to provide the installation location and the PowerProtect Data Manager server IP address.

- NOTE:** If another application agent is already installed on the client and registered to PowerProtect, ensure that you register the agent with the existing PowerProtect Data Manager server IP. When you register the agent with a PowerProtect Data Manager server that is different from the currently registered server, no warning message appears, and requests are routed to the newer server instance.
- If another application agent is already installed on the client and registered to PowerProtect, ensure that you register the agent with the existing PowerProtect Data Manager server IP. When you register the agent with a PowerProtect Data Manager server that is different from the currently registered server, no warning message appears, and requests are routed to the newer server instance.
 - If you specify a hostname or fully qualified domain name (FQDN) with an underscore (_) for the PowerProtect Data Manager server, then the communication will be done by the system's IP, if provided the system on registration.

To enable the PowerProtect Data Manager communications port that you select in the **Agent Service Port** field, ensure that the **Configure the Windows Firewall** option is selected. This option is selected by default.

When the **Configure the Windows Firewall** option is enabled, the installation creates the Windows firewall rule that allows inbound and outbound connections for the agent service process. Installation of the File System agent requires a selected port on the File System agent host (selected from port ranges 7000 to 7009 and 7012 to 7020) and the port 7001 on PowerProtect Data Manager to be open bidirectionally. These ports enable communication between the File System agent and PowerProtect Data Manager.

- NOTE:** If the Microsoft application agent is already installed and firewall rules are configured, then the **Configure the Windows Firewall** option is selected by default but disabled for the File System agent.

To select the port for communication between PowerProtect Data Manager and the File System agent, select **Agent Service Port** and specify the port number from the supported port ranges 7000 to 7009 and 7012 to 7020.

NOTE: If you do not specify a port number, the default port 7000 is used as the communications port. The ports 7010 and 7011 are used by the agent service message bus.

To select a preferred host address for communication, select **Preferred Application Host Address for Communication** and then an entry from the drop-down list. The preferred host address can be an IPv4 address, an IPv6 address, or a fully qualified domain name (FQDN).

NOTE: Selecting an FQDN ensures continued communication between PowerProtect Data Manager and the application agent host when the IP address changes. Using an FQDN is recommended for DHCP environments and other environments where static IP addresses can change.

4. Click **Install**.

The following `.msi` files are used for the installation:

- `AgentService.msi`—Installs or updates the agent service component for File System agent.
- `BBWT.msi`—Installs the block-based backups driver.
- `Fsagent.msi`—Installs the File System agent related files and folders.

5. Click **Finish**.

NOTE: If a change occurred to the PowerProtect Data Manager server IP address, the installation completes successfully but the client registration fails. To reregister the client to the correct IP address, use the **Modify** option under **Add/Remove programs** for the File System agent, and then restart the PowerProtect service agent service on the client.

Next steps

Windows installer logs are retained at `<System drive>\Users\<installing user>\AppData\Local\Temp`, and should be consulted in the event of an installation failure.

If the host is not already approved, add the file system host to the PowerProtect Data Manager server. Manage the File System agent provides more information.

Discover file system assets. Discover a file system host provides more information.

NOTE: If you change the FQDN of the client at any point, use the **Modify** option under **Add/Remove programs** to update the registration information for the File System agent, and then restart the agent service on the client to reregister the client with the PowerProtect Data Manager server.

Install the File System agent on Windows in silent mode

On Windows, review the following commands to perform a silent installation of the File System agent.

NOTE: The File System agent installer for Windows does not support the `--help` option. Running the installer program with `--help` initiates the actual installation process.

To perform the silent installation to the default path, run:

```
fsagent-19.x.x.x.exe /s PPDMServerName=<PPDM_server_IP>
```

To perform the silent installation to a different path, run:

```
fsagent-19.x.x.x.exe /s PPDMServerName=<PPDM_server_IP>  
ProductInstallPath="D:\alternate_path"
```

To enable the PowerProtect Data Manager default communications port 7000, if Windows firewall rules have not been previously configured, run:

```
fsagent-19.x.x.x.exe /s PPDMServerName=<PPDM_server_IP> EnableFirewallRules=1
```

To enable the PowerProtect Data Manager communications port 7009 and preferred application host address (as FQDN) for communications, if Windows firewall rules have not been previously configured, run:

```
fsagent-19.x.x.x.exe /s PPDMServerName=<PPDM_server_IP> EnableFirewallRules=1  
AgentServicePort=7009 PreferredAddress=bl1rV167a202.ppdn.com
```

When `EnableFirewallRules` is enabled (set to '1'), the installation creates the Windows firewall rules that allows inbound and outbound connections for the agent service process. Installation of the File System agent requires a specified port on the File System agent host (specified from the supported port ranges 7000 to 7009 and 7012 to 7020) and the port 7001 on PowerProtect Data Manager to be open bidirectionally. These ports enable communication between the File System agent and PowerProtect Data Manager.

NOTE: If you do not specify a communications port on the File System agent host, the default port 7000 is used as the communications port. The ports 7010 and 7011 are used by the agent service message bus.

The change in `EnableFirewallRules` configuration will take place only on first-time installation or upgrade of the agent service. In silent File System agent installation (in the case of coexistence), use the `ForceUpgrade=1` option to cause the silent update of common components. If you do not add this option, the silent installation fails.

NOTE: `PPDMHostName` is a mandatory option in the command line. If a value is not provided, but the agent service component has been installed by another agent, installation will succeed, but without PowerProtect, and hence it will not be possible to initiate backups from the UI. If, however, the agent service component has not been installed by another agent, then installation will fail. Specifying `ProductInstallPath` is optional, but if used, the value cannot be empty. When the `ProductInstallPath` value is provided during update or coexistence, but the install path is not the same as that of the previously installed file system or already installed agents, installation fails.

Windows installer logs are retained at `<System drive>\Users\<installing user>\AppData\Local\Temp`, and should be consulted in the event of an installation failure. In silent mode, any error message is logged only in the Windows installer logs.

Update the File System agent on Windows

If you are using an earlier version of PowerProtect Data Manager, the File System agent supports a direct update. You can update the PowerProtect Data Manager File System agent on supported Windows systems in interactive or silent mode.

Update and register the latest version of the PowerProtect Data Manager File System agent for Windows with the same PowerProtect Data Manager server in the same location.

NOTE: When you install or update the File System agent, other application agents on the system must be updated to the same version as the File System agent.

A reboot is not required after the completion of the software update.

Clustered environment requirements and considerations

The same version of the File System agent must be installed on each node in the cluster that is registered with PowerProtect Data Manager.

When the File System agent is updated from an earlier version, the following events occur:

1. Previous cluster assets with backup copies are displayed with a status of **DeLeted** in **Infrastructure > Assets > File System**.
2. Previous cluster assets are removed from protection policies, but their backup copies can be restored. These backup copies are displayed under the name of the cluster node.
3. New cluster assets are discovered, and then displayed under the name of the logical cluster host in **Infrastructure > Assets > File System**.
4. New cluster assets with the same name as previous cluster assets are automatically added to the protection policies from which the previous cluster assets were removed.

Update the File System agent on Windows in interactive mode

Use this interactive procedure to update the File System agent on supported Windows systems.


Prerequisites


- Ensure that you carry out the prerequisites provided in File System agent prerequisites.

- Download the File System agent software package.


Steps

1. In the PowerProtect Data Manager UI:


- a. Click  and then select **Downloads** from the **System Settings** menu.
- b. Select the File System agent download package for Windows, `fsagent19x_win_x64.zip`.
- c. Download the package in the location that you want to install the File System agent.

 **NOTE:** During update, or fresh installation in case of coexistence, the File System agent will be installed on previous install path of the File System agent, or the path of a previously installed Application agent.

2. Run the `fsagent-19.x.x.x.exe` program.
3. Follow the update steps in the wizard to provide the installation location and the PowerProtect Data Manager server IP address.

 **NOTE:** If another application agent is already installed on the client and registered to PowerProtect, ensure that you register the agent with the existing PowerProtect Data Manager server IP. When you register the agent with a PowerProtect Data Manager server that is different from the currently registered server, no warning message appears, and requests are routed to the newer server instance.


If PowerProtect Data Manager communications port is not enabled, you can change the firewall rule setting now as part of this update by selecting the **Configure the Windows Firewall** option.

 **NOTE:**


When the **Configure the Windows Firewall** option is enabled, the installation creates the Windows firewall rule that allows inbound and outbound connections for the agent service process. Installation of the File System agent requires a selected port on the File System agent host (selected from the supported port ranges 7000 to 7009 and 7012 to 7020) and port 7001 on PowerProtect Data Manager to be open bidirectionally. These ports enable communication between the File System agent and PowerProtect Data Manager.

If the Microsoft application agent is already installed and firewall rules are configured, then the **Configure the Windows Firewall** option is selected by default but disabled for the File System agent.

To select the port for communication between PowerProtect Data Manager and the File System agent, select **Agent Service Port** and specify the port number from the supported port ranges 7000 to 7009 and 7012 to 7020.

 **NOTE:** If you do not specify a port number, the default port 7000 is used as the communications port. The ports 7010 and 7011 are used by the agent service message bus.

To select a preferred host address for communication, select **Preferred application host address for communication** and then an entry from the drop-down list. The preferred host address can be an IPv4 address, an IPv6 address, or a fully qualified domain name (FQDN).


 **NOTE:** Selecting an FQDN ensures continued communication between PowerProtect Data Manager and the application agent host when the IP address changes. Using an FQDN is recommended for DHCP environments and other environments where static IP addresses can change. If the preferred address is an FQDN, ensure that lookup is enabled in the PowerProtect Data Manager UI. If lookup is not enabled, the preferred address defaults to an IP address.

4. Click **Upgrade**.

The following `.msi` files are used for the installation:

- `AgentService.msi`—Installs or updates the agent service component for File System agent.
- `BBBWT.msi`—Installs the block-based backups driver.
- `Fsagent.msi`—Installs the File System agent related files and folders.

5. Click **Finish**.

 **NOTE:** If a change occurred to the FQDN of the client, the installation completes successfully but the registration fails. To reregister the client to the correct FQDN, use the **Modify** option under **Add/Remove programs** for the File System agent, and then restart the agent service on the client.

Update the File System agent on Windows in silent mode

Use the following commands to perform a silent update of the File System agent on Windows.

To perform the silent update to the default path, run:

```
fsagent-19.x.x.x.exe /s PPDMSHostName=<PPDM_server_IP> ForceUpgrade=1
```

To perform the silent update to a different path, run:

```
fsagent-19.x.x.x.exe /s PPDMSHostName=<PPDM_server_IP> ForceUpgrade=1  
ProductInstallPath="D:\alternate_path"
```

To enable the default PowerProtect Data Manager communications port 7000, if Windows firewall rules have not been previously configured, run:

```
fsagent-19.x.x.x.exe /s PPDMSHostName=<PPDM_server_IP> EnableFirewallRules=1  
ForceUpgrade=1
```

To enable the PowerProtect Data Manager communications port 7009 and preferred application host address (as FQDN) for communications, if Windows firewall rules have not been previously configured, run:

```
fsagent-19.x.x.x.exe /s PPDMSHostName=<PPDM_server_IP> EnableFirewallRules=1  
AgentServicePort=7009 PreferredAddress=blrvi67a202.ppdm.com ForceUpgrade=1
```

When `EnableFirewallRules` is enabled (set to '1'), the installation creates the Windows firewall rules that allow inbound and outbound connections for the agent service process. Installation of the File System agent requires a specified port on the File System agent host (specified from the supported port ranges 7000 to 7009 and 7012 to 7020) and the port 7001 on PowerProtect Data Manager to be open bidirectionally. These ports enable communication between the File System agent and PowerProtect Data Manager.

NOTE: If you do not specify a communications port on the File System agent host, the default port 7000 is used as the communications port. The ports 7010 and 7011 are used by the agent service message bus.

The change in `EnableFirewallRules` configuration will take place only on first-time installation or upgrade of the agent service.

NOTE: `PPDMSHostName` is a mandatory option in the command line. If a value is not provided, but the agent service component has been installed by another agent, update will succeed, but without PowerProtect registration, so that it will not be possible to initiate backups from the UI. However, if the agent service component has not been installed by another agent, then update will fail. Specifying `ProductInstallPath` is optional, but if used, the value cannot be empty. When the `ProductInstallPath` value is provided during update or coexistence, but the install path is not the same as that of the previously installed file system or already installed agents, installation fails.

For File System agent silent update, or coexistence where common components are installed on the host:

- If common components are installed on the host, a File System agent update requires the additional option `ForceUpgrade=1`.
- All agents running on a client must be registered to the PowerProtect Data Manager server and must be updated to the same version.
- PowerProtect Data Manager does not support agents running different versions on the same client. Setting `ForceUpgrade=1` masks the prompt that requests users to update any other agents installed on that client to the same version to which the File System agent is being updated.

Uninstall the File System agent on Windows

On Windows, you can uninstall the File System agent with the setup file.

Steps

1. Launch `fsagent-19.x.x.x.exe`.
2. On the **Install Modification** page, select **Remove**, and then click **Next**.
3. On the **Configure Uninstallation Options** page, select **Yes** for each common component that you want to uninstall, and then click **Remove**.

4. On the **Complete the Setup** page, click **Finish**.

Results

The firewall rule created during installation for the PowerProtect Data Manager communications port is removed automatically during the uninstall operation.

Silent uninstallation commands

Steps

1. To perform a silent uninstall without uninstalling common components (such as the PowerProtect agent service or BBB), run:

```
fsagent-19.x.x.x.exe /s /uninstall
```

2. To perform a silent uninstall while also uninstalling common components, run:

```
fsagent-19.x.x.x.exe /s /uninstall UninstallPPDMAgent="1" UninstallBBBWT="1"
```

NOTE: If the File System agent is the last agent to be uninstalled, any common component that you do not uninstall remains on the host. You cannot uninstall the common components from the control panel.

Reregister the File System agent on Windows

You can use the procedure in this topic to reregister the deleted File System agent to the same PowerProtect Data Manager system.

NOTE: You can use the **Modify** option under **Add/Remove programs** for the File System agent to reregister the File System agent with PowerProtect Data Manager system only if you have not uninstalled the File System agent and the PowerProtect agent service.

If you have cleaned up the installation directories and manually uninstalled both the File System agent and the PowerProtect agent service, then you must complete the installation procedures in *Installing and updating the File System agent on Windows*.

Alternatively, you can use the `register.bat` script to register and reregister the File System agent with PowerProtect Data Manager as follows:

```
<Install_folder>\AgentService\register.bat --enable
```

Manage the File System agent

You can add a File System agent, approve and reject pending agent requests, and edit and delete existing agents.

About this task

NOTE: PowerProtect Data Manager supports the coexistence of the following agents on the same host:

- Microsoft SQL Server application agent and File System agent on Windows
- Microsoft Exchange Server application agent and File System agent on Windows
- Oracle RMAN agent and File System agent on AIX
- Oracle RMAN agent and File System agent on Linux
- SAP HANA agent and File System agent on Linux

Steps

1. Select **Infrastructure > Application Agents**.
2. In the **Application Agents** window, click **Add**.

3. Select one of the following options:

- **FGDN/IP Address**


Perform the following steps:

- a. Type the fully qualified domain name (FGDN) for the application agent.
- b. Specify the date until which the application agent is pre-approved.
- c. Click **Save**.

- **CSV Filename**

Perform the following steps:

- a. Click .

 **NOTE:** The contents of the .csv file must be in the following format, for example:

```
"ppdm.dell.com"  
"ppdm2.dell.com"
```

The **Explorer** window appears.

- b. Select the .csv file, and then click **Open**.

The file appears in the window.

- c. Select the date until which the application or File System agent is pre-approved.
- d. Click **Save**.

4. The **Auto Allow List** option is disabled by default. When **Auto Allow List** is enabled, all pre-approved Application Agents are automatically approved.


If you leave the **Auto Allow List** option disabled, select an application agent, and then select one of the following options:

- **Approve**
- **Reject**
- **Edit**, and then make the required changes.
- **Remove**

Changing the preferred host address

You can change the preferred host address that is used for communication between PowerProtect Data Manager and the application-agent host.

The preferred host address can be an IPv4 address, an IPv6 address, or a fully qualified domain name (FGDN).

 **NOTE:** Using an FGDN is only available as of version 19.14 of the application agent.

It is recommended to use an FGDN as the preferred host address in DHCP environments and other environments where static IP addresses can change. When DHCP is used instead of an IP address and the IP address of the host changes, the following benefits occur:

- Communication between PowerProtect Data Manager and the host is uninterrupted.
- Backup and restore operations succeed.
- Reregistration of the application agent is not required.

Change the preferred host address

About this task

To change the preferred host address, perform the following steps:

Steps

1. From the left navigation pane, select **Infrastructure > Application Agents**.
2. Select the entry for the host.
3. Click **More Actions** and select **Set Preferred Address**.

- From the **Preferred Address** drop-down, select the preferred address.

NOTE: DNS name resolution must be enabled to select a fully qualified domain name. If it is not enabled, enable it by selecting **Infrastructure > Application Agents** and clicking **Configure DNS Name Resolution**. After it is enabled, repeat these steps.

View application agent details


Use the **Application Agents** window in the PowerProtect Data Manager UI to monitor the registration and update status of application agents, and view details for individual application agents.

To view application agent details, from the left navigation pane, select **Infrastructure > Application Agents**.

Agent registration status displays the total number of application agents that are awaiting approval, approved, registered, or rejected.

Agent update status displays the total number of application agents that are up-to-date, available, scheduled, in progress, or failed.

NOTE: If the update of an application agent fails for any reason, the agent host is counted as available. The host is included in the total number of available applicant agents.

At the end of the **Agent update status** row, you can click  to view information about scheduled updates. The **Schedules** table appears and displays the following information:

- Update/Precheck Name
- Date and Time
- Schedule Status
- Host Count
- Actions

The lower table in the **Application Agents** window displays information about individual application agents. The following table describes the available information.

Table 6. Application agent information



Column	Description
Details	Click  in the Details column to view details and summary information for the application agent, including registration status.
Host Name	The name of the application agent host.
IP	The IPv4 or IPv6 address of the application agent host.
Registration Status	The registration status of the application agent: <ul style="list-style-type: none"> • Awaiting Approval • Pending Approval • Registered • Approved • Rejected • Expired • Accepting Certificates • Failed
OS	The operating system of the application agent host.
Agent Type	The application agent type.
Throttling Status	This column applies only to the version 19.14 or later File System agent, Microsoft SQL Server agent, and Oracle RMAN agent. The status of backup throttling on the application agent host: <ul style="list-style-type: none"> • No Throttling—CPU throttling is not set for backups on the host. • Throttling—CPU throttling is set for backups on the host.

Table 6. Application agent information (continued)

Column	Description
	<ul style="list-style-type: none"> Unsupported—CPU throttling is unsupported because the host has only pre-19.14 application agents or application agents other than the File System agent, Microsoft SQL Server agent, or Oracle RMAN agent.
CPU Throttling (hidden by default) <i>i</i> NOTE: To display this hidden column, click the icon on the lower left and then select CPU Throttling from the Show/Hide Columns list.	This column applies only to the version 19.14 or later File System agent, Microsoft SQL Server agent, and Oracle RMAN agent. The CPU utilization limit value for backup throttling on the application agent host.
Current Version	The current version of the application agent.
Update Status	The update status of the application agent host: <ul style="list-style-type: none"> Available—The PowerProtect Data Manager release is 19.14 and the application agent release is 19.10, 19.11, 19.12, or 19.13. In Progress—The update of the application agent is in progress. Up to Date—The PowerProtect Data Manager release and the application agent release are both 19.14. Scheduled—The application agent is scheduled for an update. Failed—The update of the application agent failed. Not Supported—The PowerProtect Data Manager release is 19.14 and application agent release is earlier than 19.10.

Filter and sort information

Use the filtering and sorting options to find specific application agents, and to organize the information that you see.

You can filter and sort the information that appears in table columns. Click  in the column heading to filter the information in a table column, or click a table column heading to sort that column.

Use the **Search** field to filter application agents based on a search string. When you type a keyword in the **Search** field, the PowerProtect Data Manager UI filters the results as you type. To clear the search filter, remove all keywords from the **Search** field.

Export application agent data

To export the data that is shown in the table to a .CSV file, click **Export All**.

For more information about the **Export All** functionality, see the *PowerProtect Data Manager Administrator Guide*.

Configure the centralized file-level restore service port number

By default, the centralized file-level restore (FLR) REST API service listens on TCP port 7001. This port is internal only and runs on-demand when you initiate an FLR operation, then closes after a timeout. You can change this port number to any value that does not conflict with another service.

Prerequisites

Perform this task before you start any FLR or mount operations. Verify that the new port number is not in use by any other process.

Steps

1. Create a file with the name `browsersvc.cmd` in the following location:

Operating system	Location
Windows	C:\Program Files\DPSAPPS\fsagent\settings
Linux and AIX	/opt/dpsapps/fsagent/settings

For Linux, the file owner should be root or the equivalent user for the system, with permissions set to 644.

- Using a plain text editor, open the new file.
- Add the following line to the file:

```
{ "-port": "<new-port-number>" }
```

For example, to start the FLR REST API service on port 7020: { "-port": "7020" }

- Save and close the file.

Update the application agent in the PowerProtect Data Manager UI

Learn how to perform a precheck operation and update the application agent software on one or more hosts in the PowerProtect Data Manager UI.

Prerequisites

The precheck and update operations are only available for registered clients and application agent versions 19.10 and later.

NOTE: On AIX, you cannot update the File System agent software in the PowerProtect Data Manager UI. You must use the `install.sh` script to update the File System agent on AIX from version 19.10 to 19.11.

Steps

To perform a precheck:

- From the left navigation pane in the PowerProtect Data Manager UI, select **Infrastructure > Application Agents**.
The **Application Agents** window opens.
- Select the check box next to each application agent host to be included in the precheck.
When the application agent versions on the selected hosts are 19.10 or later and the versions are earlier than the current PowerProtect Data Manager version, the **More Actions** button becomes enabled.
- Click **More Actions > Precheck Update**.
The **Precheck Update** window opens.
- On the **Schedule Precheck** page:
 - In the **Name** text box, type a name for the precheck operation.
 - Select one of the following options:
 - Precheck now**—Performs the precheck immediately.
 - Precheck later**—Schedules the precheck to occur at a later time. If you select this option, specify the date and time to perform the precheck.
 - Click **Next**.
- On the **Summary** page, review the information for the selected application agent hosts, and then click **OK**.

The precheck verifies that the application agent hosts meet the minimum update requirements, including system memory, disk space, and version requirements. If the precheck passes, PowerProtect Data Manager downloads the update software package on each application agent host.

You can monitor the progress of the precheck operation in the **System Jobs** window.

To perform an update:

- From the left navigation pane in the PowerProtect Data Manager UI, select **Infrastructure > Application Agents**.
The **Application Agents** window opens.
- Select the check box next to each application agent host to be included in the update.

NOTE: In a cluster environment, select each host of the cluster; otherwise, any unselected hosts are automatically selected for the update. It is recommended that each host of a cluster has the same application agent version.

When the application agent versions on the selected hosts are 19.10 or later and the versions are earlier than the current PowerProtect Data Manager version, the **More Actions** button becomes enabled.

8. Click **More Actions > Configure Update**.

The **Configure Update** window opens.

9. On the **Schedule Updates** page:

a. In the **Name** text box, type a name for the update operation.

b. Select one of the following options:

- **Update now**—Performs the update immediately.
- **Update later**—Schedules the update to occur at a later time. If you select this option, specify the date and time to perform the update.

c. Click **Next**.

10. On the **Summary** page, review the information for the selected application agent hosts, and then click **OK**.

On each selected host, the update performs a precheck, places the host in maintenance mode, updates the application agent, and then returns the host to normal mode.

You can monitor the progress of the update operation in the **System Jobs** window.

When the update is complete, the update status of each host changes to **Up to date** in the **Application Agents** window.

NOTE: A reboot is not required after the completion of the software update.

If the update fails:

- An error is displayed, and you must manually return the hosts to normal mode.
- Check the agent service logs for details on how to manually restore the host system.
- Check the ADM logs for more information.
- For detailed steps to downgrade to a previous version of the application agent, run the following command:

```
./pushupdate.sh -r -n
```

Moving a File System agent from one PowerProtect Data Manager system to another

To move a File System agent from one PowerProtect Data Manager system to another PowerProtect Data Manager, perform one of the following procedures:

Steps

1. Delete the asset source from File System agent, which stops the agent services.
2. On the File System agent host:
 - a. Uninstall the File System agent, which uninstalls the agent services.
 - b. Delete the `install_dir\dpsapps` folder.
 - c. Install the File System agent and provide the PowerProtect Data Manager IP address during the installation.

If you do not want to uninstall and use the existing installation and connect to the new PowerProtect Data Manager server, perform the following to move a File System agent to a new PowerProtect Data Manager system:

Steps

1. Remove the following files and folders from the File System agent host: `db`, `ssl`, `config`, and `logs`.
2. Using the **Change** option in the Windows installer or run `install.sh` on Linux or AIX and provide the IP address or FQDN of the new PowerProtect Data Manager system for registration.

Enable the File System agent after migrating File System assets to a refreshed server

A file system host is deployed using the existing data disks/volumes from the impacted host and created using the same FQDN and IP.

About this task

To migrate FS assets to a refreshed server, perform the following:

Steps

1. Delete the impacted host by deleting it from the asset source on the PowerProtect Data Manager UI.
2. Install the File System agent on the refreshed host that is created with the same FQDN and IP, and register it to the PowerProtect Data Manager.
3. Create a protection policy or add assets to an existing policy to perform backups.

Managing Storage, Assets, and Protection

Topics:

- Enable an asset source
- Discover a file system host
- Protection policies for File System agent
- Protection rules
- Cancel a File System agent backup or restore job
- Add a service-level agreement
- Extended retention for protection policies created in PowerProtect Data Manager 19.11 or earlier
- Edit the retention period for backup copies
- Delete backup copies
- Host CPU throttling
- Exclusion filters for File Systems
- Centralized restore of a file-system asset
- Updating the File System agent hostname or IP address
- Manage the PowerProtect agent service

Enable an asset source

An asset source must be enabled in PowerProtect Data Manager before you can add and register the asset source for the protection of assets.

About this task

Only the Administrator role can manage asset sources.

In some circumstances, the enabling of multiple asset sources is required. For example, a vCenter Server and a Kubernetes cluster asset source must be enabled for Tanzu Kubernetes guest cluster protection.

There are other circumstances where enabling an asset source is not required, such as the following:

- For application agents and other agents such as File System and Storage Direct, an asset source is enabled automatically when you register and approve the agent host. For example, if you have not enabled an Oracle asset source but have registered the application host through the API or the PowerProtect Data Manager user interface, PowerProtect Data Manager automatically enables the Oracle asset source.
- When you update to the latest version of PowerProtect Data Manager from an earlier release, any asset sources that were previously enabled appear in the PowerProtect Data Manager user interface. On a new deployment, however, no asset sources are enabled by default.

Steps

1. From the PowerProtect Data Manager user interface, select **Infrastructure** > **Asset Sources**, and then click + to reveal the **New Asset Source** tab.
2. In the pane for the asset source that you want to add, click **Enable Source**.
The **Asset Sources** window updates to display a tab for the new asset source.

Results

You can now add or approve the asset source for use in PowerProtect Data Manager. For a vCenter server, Kubernetes cluster, SMIS Server, or PowerProtect Cloud Snapshot Manager tenant, select the appropriate tab in this window and click **Add**. For an application host, select **Infrastructure** > **Application Agents** and click **Add** or **Approve** as required.

NOTE: Although you can add a Cloud Snapshot Manager tenant to PowerProtect Data Manager in order to view its health, alerts, and the status of its protection, recovery, and system jobs, you cannot manage the protection of its assets from

PowerProtect Data Manager. To manage the protection of its assets, use Cloud Snapshot Manager. For more information, see the *PowerProtect Cloud Snapshot Manager Online Help*.

Disable an asset source

If you enabled an asset source that you no longer require, and the host has not been registered in PowerProtect Data Manager, perform the following steps to disable the asset source.

About this task

NOTE: An asset source cannot be disabled when one or more sources are still registered or there are backup copies of the source assets. For example, if you registered a vCenter server and created policy backups for the vCenter Server virtual machines, then you cannot disable the vCenter Server asset source. But if you register a vCenter server and then delete it without creating any backups, you can disable the asset source.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Asset Sources**, and then select the tab of the asset source that you want to disable.
If no host registration is detected, a red **Disable** button appears.
2. Click **Disable**.

Results

PowerProtect Data Manager removes the tab for this asset source.

Delete an asset source

If you want to remove an asset source that you no longer require, perform the following steps to delete the asset source in the PowerProtect Data Manager UI.

About this task

Only the Administrator role can manage the asset sources.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Asset Sources**, and then select the tab for the type of asset source that you want to delete.
2. Select the asset source name in the asset source list, and then click **Delete**.
3. At the warning prompt that appears, click **Continue**.
The asset source is deleted from the list.

Results

PowerProtect Data Manager removes the specified asset source in the **Asset Sources** window.

Any associated assets that are protected by the protection policy are removed from the protection policy and their status is changed to `deleted`. These assets are removed automatically as part of daily PowerProtect Data Manager cleanup after all associated backup copies have been deleted. These assets can also be removed manually. The *PowerProtect Data Manager Administrator Guide* provides details on how to remove assets from PowerProtect Data Manager.

The copies of assets from the asset source are retained (not deleted). You can delete the copies from the copies page, if required.

Discover a file system host

Perform the following steps to discover a file system host as an asset source in the PowerProtect Data Manager UI.

Steps

1. Select **Infrastructure > Asset Sources**.

The **Asset Sources** window appears.

2. Select the agent host tab.
3. Select the file system host and click **Discover**.


The **Initiate Discovery** dialog appears with an option to immediately start a full discovery of the assets on the host.


NOTE:

- From the agent host tab, you can click **Discover** at any time if any additions or other changes to your asset sources have taken place outside of the PowerProtect Data Manager environment. Asset discovery is also initiated by default after registration of the host to PowerProtect Data Manager and at hourly intervals. Discovery time is based on networking bandwidth. Each time you initiate a discovery process, the resources that are discovered and those that are handling the discovery impact system performance.
- Use a unique combination of a hostname (FGDN) and asset name while discovering assets. Asset discovery fails when the same hostname (FGDN) is used for more than one agent.

4. Click **Yes**.

Results

-  **NOTE:** From PowerProtect Data Manager 19.13, all the Linux mount points such as ext3, ext4, XFS, and BTRFS can be discovered and protected even if they are not listed in the `/etc/Fstab` entry.

When the File System asset source is configured correctly, you can add the file system assets to a PowerProtect Data Manager protection policy. Go to **Infrastructure > Assets**, and then select the **File Systems** tab. Use  to switch between a list view of all file system assets for all of the discovered File System hosts and a hierarchical view of assets within each File System host that has been discovered in PowerProtect Data Manager.

Protection policies for File System agent

Protection policies define sets of objectives that apply to specific periods of time. These objectives drive configuration, active protection, and copy-data-management operations that satisfy the business requirements for the specified data. Each policy type has its own set of user objectives.

Only the Administrator role can create or edit protection policies.

The next topics discuss protection policy terminology, concepts, and available options for planning protection policies. All protection policies require a primary backup or retention objective. The remaining objectives are optional.

Supported protection policy purposes

These points provide generalized descriptions of PowerProtect Data Manager behavior. The different asset types prescribe asset-specific actions and conditions for each purpose:

- **Centralized Protection**—PowerProtect Data Manager manages the complete protection life cycle. The backup trigger is part of the protection policy.
- **Self-Service Protection**—The local application on the host handles backing up data and provides backup triggers. The local application passes backup data to PowerProtect Data Manager, which manages the rest of the protection life cycle.
- **Exclusion**—Marks the specified assets as intentionally exempt from data protection operations. Otherwise, assets that are not assigned to any protection policy are reported as unprotected. PowerProtect Data Manager restricts the specified assets from being assigned to other protection policies.

Supported protection policy objectives

Each objective of a protection policy accomplishes one of the following purposes:

- **Primary Backup**—Creates and updates the catalog of backups for data that is protected by a centralized protection policy.
- **Primary Retention**—Creates and updates the catalog of backups for data that is protected by a self-service protection policy. PowerProtect Data Manager receives the backup data from a local application on the asset.
- **Replication**—Sends a replica of designated backups or retentions, along with the associated metadata, to a remote server for additional redundancy.
- **Extended Retention**—Extends the retention period for designated backups or retentions for long-term purposes. This type of objective is only supported for protection policies that were created in PowerProtect Data Manager 19.11 or earlier.
- **Cloud Tier**—Moves the designated backups, retentions, or replicas from the local protection storage to an associated cloud unit for long-term storage or archival purposes.

Primary backup objective

A primary backup objective contains a backup target and one or more schedules.

The backup target determines where and how PowerProtect Data Manager stores the asset backups:

- Which protection storage system and storage unit to use.
- Which network interface is necessary to reach the selected protection storage system.
- Whether to enable retention locking that protects against accidental deletion or tampering.

The schedules determine what and when PowerProtect Data Manager backs up:

- The types of backups to perform.
- When backups should start.
- How often to perform each type of backup.
- How long to retain each type of backup.
- The times when PowerProtect Data Manager can and cannot perform backups.

Primary retention objective

A primary retention objective contains a retention target and retention rules.

The retention target determines where and how PowerProtect Data Manager stores the asset backups:

- Which protection storage system and storage unit to use.
- Which network interface is necessary to reach the selected protection storage system.
- Whether to enable retention locking that protects against accidental deletion or tampering.

The retention rules determine how PowerProtect Data Manager handles backup data from the asset:

- How long to retain each type of backup.
- Whether different types of backup should be retained for the same period.
- Whether to extend the retention period for specific backups.

Replication objective

A replication objective contains a replication target and one or more schedules.

The replication target determines where and how PowerProtect Data Manager stores the replicas:

- Which protection storage system and storage unit to use.
- Which network interface is necessary to reach the selected protection storage system.
- Whether to enable retention locking that protects against accidental deletion or tampering.

The schedules determine what and when PowerProtect Data Manager replicates:

- The types of backups to replicate.
- When replication should start.
- How often to replicate.
- How long to retain each replica.

- The times when PowerProtect Data Manager can and cannot replicate.

Extended retention objective

Protection policies that were created in PowerProtect Data Manager 19.9 or earlier might have extended retention objectives. However, you cannot edit or add new extended retention objectives.

Cloud tier objective

A cloud tier objective contains tiering rules:

- Which backups, retentions, or replicas to move to the cloud tier.
- When to move data to the cloud tier.

File System replication triggers

PowerProtect Data Manager orchestrates protection policy replication objectives independently of the primary backup. When you add a replication objective to a policy, select one of the available triggers.

The default replication trigger is a schedule window that you define by setting a recurrence period plus start and end times. Replication occurs during the defined window. For example, every day between 8 p.m. and 12 a.m.

You can also trigger replication immediately after the completion of the associated primary backup, whether scheduled or manual. At the start of the primary backup, PowerProtect Data Manager generates an associated replication job that remains queued until the end of the protection job. If the backup fails or completes with exception, the associated replication job is skipped. Restarting the protection job queues the associated replication job again.

When you create a replication objective, you can specify either scheduled replication or replication after backup completion, which is applicable to both centralized and self-service protection policies.

NOTE: For replication after backup completion, it is recommended that you update the application agents to the latest version. Depending on the type of backup, the following versions are required to ensure that replication occurs immediately after the backups complete:

- For self-service primary backups, update all application agents to PowerProtect Data Manager version 19.12 or later.
- For centralized primary backups, update all application agents to PowerProtect Data Manager version 19.11 or later.

If you want to replicate only specific backups, perform a manual replication of these backups in advance.

Using a schedule can help you manage network traffic by replicating during off-peak hours. However, for larger backup sets, the primary backup may not finish before the start of the replication schedule, which creates a replication backlog. Replication after backup completion prevents a replication backlog from forming.

To prevent data loss, the replication after backup completion trigger replicates new backups from the primary objective and any outstanding backups that have not yet replicated.

Roadmap for planning a protection policy

Before you create a protection policy, assemble the following information:

1. Identify a supported purpose for this protection policy that corresponds to your goal.
For example, performing centralized protection or application aware backups.
2. Identify what supported objectives this protection policy should accomplish. Supported protection policy objectives provides more information.
For example, primary backup and replication to a remote server.
3. Use the following tables to list all required primary backup or retention types and their associated retention periods. Include any additional full backups for extended retention.
4. For centralized, crash consistent, or application aware protection, list the backup frequency and start/end times for each primary backup type.
5. Identify a protection storage system and storage unit where PowerProtect Data Manager should store the primary backup or retention.

For this target, identify any required virtual networks or interfaces, and whether retention locking is required.

6. If replication is required, identify the remote protection storage system and storage unit.

Table 7. Required backup and retention types

Primary backup or retention types	Retention period (with units)	Retention period units
Example (full)	1 week	weeks
Full		
Synthetic full		
Other:		
Other:		
Other:		

Table 8. Centralized protection schedules

Primary backup types	Creation frequency (with units)	Start and end
Example (full)	1 week	08:00 PM to 08:00 AM
Full		
Synthetic full		
Other:		
Other:		
Other:		

Table 9. Additional full backups for extended retention

Primary backup or retention types	Created every (week/month/year)	On the (day of week/month/year)	Retention period
Retained full backup 1 (example)	Week	Saturday	52 weeks
Retained full backup 2:			
Retained full backup 3			

Before you add a protection policy for file system protection

Review the limitations, prerequisites, and best practices in this section and for enabling the File System agent before you continue.

Additional procedures for system administrators

The *PowerProtect Data Manager Administrator Guide* provides more information about procedures that require the Administrator role. The PowerProtect Data Manager system administrator typically performs these procedures, some of which have server-wide effect.

Many of these procedures are not specific to this asset type. However, some procedures in the *PowerProtect Data Manager Administrator Guide* are prerequisites for, or also applicable to, asset protection. For example:

- Adding and configuring protection storage, including storage units
- Adding and configuring virtual networks
- Managing protection policies, including adding or removing assets, and disabling a policy, and protection rules
- Managing backups, such as editing retention periods and deleting backup copies

- Managing any running jobs

Review these procedures when required for your environment or when called out by a task in this guide. If required, coordinate with your PowerProtect Data Manager system administrator or backup and restore administrators.

Maintaining protection policies

Refer to the *PowerProtect Data Manager Administrator Guide* for more information about editing, disabling, or deleting protection policies.

You can change any of the following information for an existing enabled or disabled protection policy:

- Policy name and description
- Adding or removing assets from the policy
- Backup and replication schedule
- Backup optimization mode
- Settings for network interface, storage target, storage unit, and retention lock.

You cannot modify a protection policy type or purpose. For these actions, add a policy with the new type or purpose. Storage quotas cannot be changed by editing a policy.

NOTE:

Once you save changes for an enabled or disabled policy, most changes take effect immediately. For a disabled policy's primary backup schedules, however, the changes do not take effect until you reenable the policy, since these schedules do not run in **Disabled** state.

Protection policy limitations

Observe the following information when planning data protection:

Exclusion filters

- When an exclusion filter is applied to a protection policy, the File System agent performs file-based backups of the protected assets. File-based backups traverse through the entire directory structure of the file system to back up all the files in each directory of the file system. While file-based backups can provide additional capabilities such as exclusion, these backups take longer to complete when compared to block-based backups.
- Exclusion filters cannot be applied to self-service protection policies or to backups taken through the self-service CLI.

Protection policy prerequisites

Before you configure a protection policy for asset protection, observe the following points and perform the following actions:

Notes

- You can only protect an asset with one policy at a time. Assets can move between protection policies, depending on the protection rule priorities. Protection rules do not automatically move assets that were manually added to a policy to a different policy.
- Before scheduling weekly, monthly, or yearly backups, set the PowerProtect Data Manager time zone to the local time zone. Otherwise, the backup still runs but is triggered based on the PowerProtect Data Manager time zone.
- Time-based exclusion filters require an NTP server to synchronize the time on the PowerProtect Data Manager and File System assets.

Windows, Linux, and AIX

- Ensure that your host is supported. The E-Lab Navigator provides software compatibility information for PowerProtect Data Manager.
- Changing the retention periods for specific backup types requires File System agent 19.9 or later.

- Upon requesting a backup (file-based or block-based), the status of the protection policy becomes **Queued**. This status switches to **Running** only after the system begins writing the backup to protection storage.

Initial configuration

- Enable the File System agent asset source.
- Register the application hosts with PowerProtect Data Manager.
- Perform a discovery of the application hosts.
- Where applicable, configure any necessary Service Level Agreements (SLAs). The *PowerProtect Data Manager Administrator Guide* provides instructions.
- (Optional) Configure the file system parallel backup setting.
- (Optional) Configure asset multi-streaming for file-based backups.

Objectives

- For replication after backup completion, PowerProtect Data Manager 19.12 or later and application agents 19.12 or later are required. It is recommended that you update the application agents to the latest version.
- To move a backup or replica to Cloud Tier, the corresponding objectives must have a retention time of 14 days or more.

Storage

- Add protection storage.

The *PowerProtect Data Manager Administrator Guide* provides more information about working with storage units, such as the relationships between storage units and policies, and applicable limitations.

- Before you add a replication objective, add remote protection storage for the replication target.
- Before you add a Cloud Tier objective, PowerProtect Data Manager requires the discovery of protection storage with a configured Cloud unit.

NOTE: PowerProtect Data Manager does not support the automatic retention lock (ARL) setting on the DD system. The option to create a storage unit during protection policy configuration does not support compliance mode retention locking, only governance mode. To use compliance mode retention locking, create and configure a storage unit before you configure an associated protection policy. If you enable retention locking and select a storage unit where the retention lock mode is *None*, the retention lock defaults to governance mode. The *PowerProtect Data Manager Administrator Guide* provides more information.

Networking

If applicable, complete all of the virtual network configuration tasks before you assign any virtual networks to the protection policy. The *PowerProtect Data Manager Administrator Guide* provides more information.

Enable file-level restores

Steps

1. Log in to the system you are restoring from as the root user.
2. Install the iSCSI client packages.
The operating system documentation provides package installation procedures.
3. For **Service Start**, choose **Manually**.
4. Click **OK**.

Protection policy best practices

Before you create a protection policy, note the following best practices:

- In the **Protection Policies** window, you can export protection policy data by using the **Export All** functionality.

Understanding backup terminology and managing backup frequency

When scheduling backups in a protection policy, be aware of the following:

- Different protection-policy types can use different terminology to describe available backup levels. This terminology can differ not only between protection-policy types, but also from traditional terminology.
- To avoid high CPU usage that can lead to failure issues, do not schedule backups more often than recommended.

To understand the different backup levels to manage backup frequencies, see the following table.

Table 10. Backup terminology and frequency

Protection-policy types	Available backup levels	Description	Equivalent traditional terminology	Recommended minimum backup interval
File System centralized	Full	All the data is backed up.	Full	Monthly
	Synthetic Full	Only the data that has changed since the last synthetic-full or full backup is backed up. An operation to merge these changes with the last synthetic-full or full backup produces a full backup in storage. Only the changed blocks are copied over the network, but the result is still a full backup in storage.	A differential backup is performed, followed by a merge operation that produces a full backup in storage.	12 hours

NOTE: In some situations, a full backup might be performed even though a synthetic-full backup was scheduled. Possible reasons for a full backup include the following:

- There is no existing full backup.
- The size of a volume has changed.
- There has been a file path change.
- The asset host has been rebooted.

The backup frequency of log, differential, incremental-cumulative, incremental-differential, and incremental backups cannot be greater than the backup frequency of either full or synthetic-full backups. If you attempt to add or edit a protection policy that uses an invalid backup frequency, PowerProtect Data Manager prevents you from saving the protection policy. You can increase the backup frequency of a protection policy by scheduling more full or synthetic-full backups with different retention times to meet your requirements.

Best practices for file system backups

Consider the following best practices for file system backups.

- Ensure that subsequent backup schedules do not overlap.

If a full backup is in progress when the next incremental backup starts, then the incremental backup is promoted to a full backup. If the schedules overlap, the incremental backup continues to get promoted to a full backup because a full backup has not completed. Also, overlapping backup schedules might fail as a result of concurrent snapshots or other limitations that are caused by the underlying system.

To identify the optimal backup window between full backups and the first incremental backup, measure the backup time that is required for a small asset. Use this time as an indicator to assess the backup window for assets that are larger in size. For example:

- Data backup speed = asset size / time taken to back up the asset
 - Total backup size = sum of the total size of the assets to be backed up for a full backup
 - Maximum backup window = total backup size / data backup speed
 - Optimal backup window = maximum backup window / parallelism value
- If the data source is encrypted or compressed, the DD system provides limited deduplication. This increases the time that is taken to back up the data, as compared to data that is not encrypted or compressed.

Add a protection policy for file system protection

The following tasks complete the **Add Policy** wizard, which is specific to each asset type:

1. Select a protection policy type
2. Select a protection policy purpose
3. Select protection policy assets
4. Configure file exclusions
5. Add protection policy objectives
6. Configure protection policy options
7. Review the protection policy summary

After you have added a protection policy, Viewing protection policy jobs provides more information about ongoing protection activities that correspond to each objective.

Select a protection policy type

Open the **Add Policy** wizard and select the protection policy type.

Steps

1. From the left navigation pane, select **Protection > Protection Policies**.
The **Protection Policies** window appears.
2. In the **Protection Policies** window, click **Add**.
The **Add Policy** wizard appears.
3. On the **Type** page, specify the following fields, and then click **Next**:
 - **Name**—Type a descriptive name for the protection policy.
 - **Description**—Type a description for the policy.
 - **Type**—Select **File System**.

Results

The wizard moves to the **Purpose** page. Continue to Select a protection policy purpose.

Select a protection policy purpose

Select from the list of supported purposes.

Steps

On the **Purpose** page, select one of the following options to indicate the purpose of the new protection policy group, and then click **Next**:

- **Centralized Protection**—PowerProtect Data Manager manages the complete protection life cycle. The backup trigger is part of the protection policy.
- **Self-Service Protection**—The local application on the host handles backing up data and provides backup triggers. The local application passes backup data to PowerProtect Data Manager, which manages the rest of the protection life cycle.
- **Exclusion**—Marks the specified assets as intentionally exempt from data protection operations. Otherwise, assets that are not assigned to any protection policy are reported as unprotected. PowerProtect Data Manager restricts the specified assets from being assigned to other protection policies.

Results




The wizard moves to the **Assets** page. Continue to Select protection policy assets.

Select protection policy assets


Steps

1. On the **Assets** page, select the unprotected assets that you want to add to the backup of this protection policy group.

The window enables you to filter by asset name to locate the required assets.

You can use    to switch between a list view of all assets discovered by PowerProtect Data Manager and a hierarchical view to display the assets in a tree structure underneath each host. A hierarchical view can be helpful if you have added multiple file systems and need to more easily identify which assets belong to which host. You can also select **Show only unprotected assets** to further filter the list.

NOTE:

- When you select the entire set of an asset including the server/host or when selecting only the server/host of an asset, an icon,  appears and a dynamic protection rule is created to ensure that the selected server/host are protected by the protection policy.
- The protection rules that are created while adding or editing the protection policy, cannot be edited or deleted from the **Protection Rules** window.

2. Click **Next**.

Results

For exclusion policies, the wizard moves to the **Summary** page. Continue to Review the protection policy summary.

For centralized protection policies, the wizard moves to the **File Exclusions** page. Continue to Configure file exclusions.

For self-service protection policies, the wizard moves to the **Objectives** page. Continue to Add protection policy objectives.

Configure file exclusions

You can specify properties that cause PowerProtect Data Manager to exclude matching discovered files or folders from backups. For example, swap files. You can define a new exclusion filter on the **File Exclusions** page or apply a saved filter.

About this task

You can add up to four exclusion filters. When you add multiple conditions, a file is excluded only if it meets all filter conditions. Within a filter, you can add a condition only once.

Exclusion filters for File Systems provides more information.

Steps

1. To enable exclusions, set **Would you like to enable exclusions?** to **Enabled**.
The wizard displays the available filtering conditions and the option to use an existing filter or group of filters as a template.
2. To apply or import a saved filter:
 - a. Click **Add Saved Filter**.
The **File Exclusion Filter Groups** page appears and displays the existing saved filters and filter groups.
 - b. Use the **Search name of a Filter Group** field to filter the list.
 - c. For any saved filter or filter group, select **USE AS TEMPLATE** or **USE THIS GROUP**.
Both options transfer the selected filter or filter group to the **File Exclusions** page. However, the **USE AS TEMPLATE** option enables you to edit the filtering conditions from the saved filter. The **USE THIS GROUP** option treats the saved filter as an immutable preset.
 - d. When you are finished applying or importing filters, click **Finish**.
3. To define a new filter:
 - a. Select one or more filtering conditions.

+File Size	Exclude files and folders that are larger, smaller, or equal to a specified size. Select Equals , Greater than , or Less than and then specify a size value with units.
+File Type	Exclude files or folders, based on file type. Specify a file name extension or multiple file name extensions. Separate multiple values with commas.
+Modified Time	Exclude files or folders that were modified before or after a certain date. Specify a date in either the After or Before field. For NAS shares, provide the modified time in Coordinated Universal Time (UTC). If you do not convert to UTC, the filter might exclude incorrect files or folders.
+Folder Path	Exclude files and folders in a specific path. Specify the file path, and then enclose the file path in quotations. You can specify an absolute or relative path.

- b. When you are finished building the filter, click **Add Filters**.
The wizard displays the new exclusion filter as a group.
 - c. Type a name and description for the filter and then click **Save**.
4. When you are finished adding exclusion filters, click **Next**.

Results

The wizard moves to the **Objectives** page. Continue to Add protection policy objectives.

Add protection policy objectives

Use these instructions to add the objectives that you identified during policy planning.

- Configure a primary backup objective
- Configure a primary retention objective
- Configure a replication objective
- Configure a cloud tiering objective

Optionally, select a policy-level Service Level Agreement (SLA) from the **Set Policy Level SLA** list, or select **Add** to open the **Add Service Level Agreement** wizard and create a policy-level SLA.

The *PowerProtect Data Manager Administrator Guide* provides instructions.

Configure a primary backup objective

Primary backup objectives apply to all protection policy purposes except self-service and exclusion policies.

Prerequisites


NOTE:

When a new asset is added to a protection policy during a scheduled backup window, the backup starts right away. However, if an asset is added to a protection policy outside of the scheduled backup window, the backup does not start until the next time that backups are configured to run.

If a new asset is added to a protection policy that has a weekly or monthly backup schedule and the current time is within the scheduled **Start** and **End** times, the backup runs right away, regardless of the date. If the current time is not within the scheduled **Start** and **End** times, the backup does not start until the next time that backups are configured to run.

-  NOTE: Any backup that starts before the **End** time continues until completion.

About this task

 CAUTION: The retention period of synthetic full backups must be less than or equal to the retention period of full backup copies. If you set a shorter retention period for a synthetic full backup than for the corresponding full backup, then data loss might occur and you might be unable to recover the point-in-time copies.

By default, the retention period for the full backup is the same as that for the synthetic full backup. You can, however, specify a retention period for the full backup that is longer than the retention period for the synthetic full backup.

Steps

1. Click **Add** under **Primary Backup**.
The **Add Primary Backup** dialog appears.
2. On the **Target** pane, specify the following fields:
 - a. **Storage Name**—Select a backup destination from the list of protection storage, or select **Add** to add protection storage and complete the details in the **Storage Target** dialog.
 - b. **Storage Unit**—Select whether this protection policy should use a **New** storage unit on the selected protection storage system, or select an existing storage unit from the list.

Hover over a storage unit to view the full name and statistics for available capacity and total capacity, for example, `testvmpolicy-daily-123ab (300 GB/1 TB)`.

The **Space** field indicates the total amount of space, and the percentage of available space, on the protection storage system.

When you select **New**, PowerProtect Data Manager creates a storage unit on the selected protection storage system upon policy completion. The storage unit name follows the format **policy name-hostname-unique identifier**.

For example, **testvspolicy-daily-123cd**.

- c. **Network Interface**—Select a network interface from the list, if applicable.
 - d. **Retention Lock**—Move the **Retention Lock** slider to the right to enable retention locking for these backups.
The retention lock mode setting comes from the configuration of the selected storage unit. When you enable retention locking, the **Retention Lock Mode** field displays the corresponding storage unit setting.
Setting a retention lock applies to the current backup copy only, and does not impact the retention lock setting for existing backup copies.
 - e. **SLA**—From the list, select an existing service level agreement that you want to apply to this objective, or select **Add** to create an SLA within the **Add Backup Service Level Agreement** wizard.
The *PowerProtect Data Manager Administrator Guide* provides instructions.
3. On the **Schedules** pane, specify the following fields to schedule a synthetic full backup of this protection policy:
 - a. **Create a Synthetic Full backup every**—Specify how often to create a synthetic full backup. A synthetic full backs up only the data that has changed since the last backup, to create a new, full backup.
 - b. **Retain for**—Specify the retention period for the synthetic full backup.
 - c. **Start and End**—The activity window. Specify a time of day to start the synthetic full backup, and a time of day after which backups cannot be started.
 - d. Click **Save**.
 4. To periodically force a full (level 0) backup, click **Add backup** and then specify the following fields to schedule the full backup:

NOTE: When you force a full backup, the backup chain resets.

- a. **Create a Full backup every**—Specify whether you want to create an hourly, daily, weekly, monthly, or yearly full backup.
 - b. **Repeat on**—Depending on the frequency of the full backup schedule, specify the hour of the day, day of the week, or the date of the month that the full backup will occur.
 - c. **Retain for**—Specify the retention period for this full backup.
 - d. **Start and End**—The activity window. Specify a time of day to start the full backup, and a time of day after which backups cannot be started.
 - e. Click **Save**.
5. Click **Add Backup** and repeat the procedure for creating full backups if you want to create additional backup copies at different intervals with different retention periods.

Within this protection policy, when a full schedule conflicts with another full backup schedule, a message appears, indicating that there is a conflict. Schedule occurrences can conflict with each other when the activity windows are identical or occur entirely within the same time range. To avoid full schedule conflicts in a policy, edit the activity windows.

If you proceed with conflicting schedules, the backup of the lower priority schedule will be skipped. Schedule priority is ranked according to the following criteria:

- Full schedules have a higher priority than Synthetic Full schedules.
- For schedules of the same backup type, the schedules that run less frequently have a higher priority than schedules that run more frequently.
- For schedules with the same backup type and frequency, the schedule with the longest activity window has the higher priority. If the activity windows are also identical, only one of these schedules will run.

When a schedule conflict between full backups occurs, PowerProtect Data Manager runs the full backup with the longest retention period.

6. Click **Save** to save the changes and return to the **Objectives** page.
The **Objectives** page updates to display the name and location of the storage target under **Primary Backup**.

Next steps

After completing the objective, you can change any details by clicking **Edit** next to the objective.

Configure any remaining objectives. When you have configured all required objectives, click **Next**.

The wizard moves to the **Options** page. Continue to Configure protection policy options.

Configure a primary retention objective

Retention objectives apply to self-service protection policies.

About this task

If you only want to replicate specific backups, perform a manual replication of these backups in advance.

Steps

1. Click **Add** under **Primary Retention**.
The **Add Primary Retention** dialog appears.
2. On the **Target** pane, specify the following fields:
 - a. **Storage Name**—Select a backup destination from the list of protection storage, or select **Add** to add protection storage and complete the details in the **Storage Target** dialog.
 - b. **Storage Unit**—Select whether this protection policy should use a **New** storage unit on the selected protection storage system, or select an existing storage unit from the list.

Hover over a storage unit to view the full name and statistics for available capacity and total capacity, for example, **testvmpolicy-daily-123ab (300 GB/1 TB)**.

The **Space** field indicates the total amount of space, and the percentage of available space, on the protection storage system.

When you select **New**, PowerProtect Data Manager creates a storage unit on the selected protection storage system upon policy completion. The storage unit name follows the format **policy name-hostname-unique identifier**.
For example, **testvmpolicy-daily-123cd**.
 - c. **Network Interface**—Select a network interface from the list, if applicable.
 - d. **Retention Lock**—Move the **Retention Lock** slider to the right to enable retention locking for these backups.
The retention lock mode setting comes from the configuration of the selected storage unit. When you enable retention locking, the **Retention Lock Mode** field displays the corresponding storage unit setting.

Setting a retention lock applies to the current backup copy only, and does not impact the retention lock setting for existing backup copies.
 - e. **SLA**—From the list, select an existing service level agreement that you want to apply to this objective, or select **Add** to create an SLA within the **Add Service Level Agreement** wizard.
The *PowerProtect Data Manager Administrator Guide* provides instructions.
3. On the **Retention (Self Service)** pane, change any required retention periods.
By default, all backup types share the same retention period.
4. To change the retention periods for specific backup types:
 - a. Clear **Set the same retention time for all backup types**.
 - b. Change the **Retain <backup_type> For** field values as required.

After changing this option, you can create additional backup patterns with different retention periods. For example, you can add a full backup pattern **Retain full backups created every week on the Monday and Tuesday for 2 months**.

Self-service retentions created with older versions of the File System agent continue to use the same retention period for full and synthetic full backups.
5. Click **Save** to save the changes and return to the **Objectives** page.
The **Objectives** page updates to display the name and location of the storage target under **Primary Retention**.

Next steps

After completing the objective, you can change any details by clicking **Edit** next to the objective.

Configure any remaining objectives. When you have configured all required objectives, click **Next**.

The wizard moves to the **Options** page. Continue to Configure protection policy options.

Configure a replication objective

Optionally, replicate the primary backup or retention to a remote server for added protection. You can specify either scheduled replication or replication after backup completion.

Prerequisites

- NOTE:** When creating multiple replicas for the same protection policy, it is recommended to select a different storage system for each copy.
- NOTE:** If you select a storage unit that is the target of another objective for the same policy, the UI issues a warning. The *PowerProtect Data Manager Administrator Guide* provides information about replicating to shared protection storage to support PowerProtect Cyber Recovery. Verify the storage targets and the use case before you continue.

About this task

For replicas of centralized backups, when you set retention periods for different backup types, any undefined types use the full backup retention period. For example, if you do not define a log backup in the primary objective, the log backup for the replication objective is also undefined. After you run a manual log backup, replicas of that log backup use the same retention period as the full backup.

Steps

1. Next to the primary backup or retention objective, click **Replicate**.
An entry for **Replicate** appears to the right of the primary backup or retention objective.

2. Under **Replicate**, click **Add**.

The **Add Replication** dialog appears, with information in the left pane for each schedule that has been added for the primary objective of this protection policy.

- NOTE:** PowerProtect Data Manager replicates backups for all the listed schedules. You cannot select individual schedules for replication.
3. Select a storage target:
 - a. **Storage Name**—Select a replication destination from the list of protection storage, or select **Add** to add protection storage and complete the details in the **Storage Target** window.
 - b. **Storage Unit**—Select whether this protection policy should replicate to a **New** storage unit on the selected protection storage system, or select an existing storage unit from the list.
 - c. **Network Interface**—Select a network interface from the list, if applicable.
 - d. **Retention Lock**—Move the **Retention Lock** slider to the right to enable retention locking for these replicas.
The retention lock mode setting comes from the configuration of the selected storage unit. When you enable retention locking, the **Retention Lock Mode** field displays the corresponding storage unit setting.
 - e. **SLA**—Select an existing replication service level agreement that you want to apply to this schedule from the list.
Or, select **Add** to create a replication SLA within the **Add Replication Service Level Agreement** wizard.

The *PowerProtect Data Manager Administrator Guide* provides more information about replication targets.

4. Select when to replicate the backups or retentions:

File System replication triggers provides more information.

- To replicate after the backup finishes, move the **Replicate immediately upon backup completion** slider to on.
- For scheduled replication, move the **Replicate immediately upon backup completion** slider to off, and then complete the schedule details in the **Add Replication** dialog.

For replication of the primary backup, the schedule frequency can be every day, week, month, or x hours.

For daily, weekly, and monthly schedules, the numeric value cannot be modified. For hourly, however, you can edit the numeric value. For example, if you set **Create a Full backup every 4 hours**, you can set a value of anywhere from 1 to 12 hours.

By default, all replicas of the primary objective inherit the retention period from the **Retain For** value of the synthetic full and full backup schedules.

5. To specify a different retention period for replicas of different backup or retention types:

CAUTION: If you set a shorter retention period for the replicas of additional backup types than for the corresponding full backup, you may be unable to recover from those replicas. The additional backup types include log, incremental, differential, and so on, where applicable.

- a. Clear **Set the same retention time for all replicated copies**.
 - b. Click **Edit** in the row of each schedule that you want to change.
 - c. Update the value in the **Retain For** field.
 - d. Click **Save**.
6. Click **Save** to save your changes and return to the **Objectives** page.

Next steps

Configure any remaining objectives. When you have configured all required objectives, click **Next**.

The wizard moves to the **Options** page. Continue to Configure protection policy options.

Configure a cloud tiering objective

For some protection policy types, you can add a Cloud Tier objective to a protection policy to move local full backups to the Cloud Tier after a predefined number of days.

Prerequisites

To move a backup or replica to Cloud Tier, the corresponding objectives must have a retention time of 14 days or more.

Cloud Tiering happens at 00:00 UTC each day. Depending on your time zone, this time may be within business hours and thus Cloud Tiering may impact available network bandwidth. Cloud Tiering applies to both centralized and self-service protection policies.

About this task

The *PowerProtect Data Manager Administrator Guide* provides more information about adding a Cloud objective and managing or working with Cloud Tier backups and replicas.

Steps

1. Click **Cloud Tier** next to or under one of the following:
 - A primary backup objective.
 - A primary retention objective.
 - A replication objective.The wizard creates an entry for **Cloud Tier** to the right of, or below, the associated objective.
2. Under the entry for **Cloud Tier**, click **Add**.

The **Add Cloud Tier Backup** dialog appears, with summary information for the parent objective to indicate whether you are adding this Cloud Tier objective for the primary objective or the replication objective.
3. To tier the backups from all the full primary backup or replication schedules of this policy, keep the **All applicable full backups** slider to the right.

Otherwise, move the slider to the left and select one or more full schedules to tier.
4. Set the following parameters:
 - Select the appropriate Cloud Unit from the **Cloud Target** list.
 - For **Tier After**, set a time of 14 days or more.

NOTE: If either of the following conditions are true, you can still select this schedule for tiering:

 - The retention period of a schedule is less than the minimum 14 days required before tiering occurs.
 - The retention period of a schedule is less than the value in the **Tier After** field.

However, you must edit the retention period of this schedule, or its backup or replica, to a value greater than the **Tier After** field before the retention period of the copy expires. Otherwise, PowerProtect Data Manager will not move the backup or replica of this schedule to the cloud tier.
5. Click **Save** to save your changes and return to the **Objectives** page.

Next steps

Configure any remaining objectives. When you have configured all required objectives, click **Next**.

The wizard moves to the **Options** page. Continue to Configure protection policy options.

Configure protection policy options

On the **Options** page, select any additional options that are required for the policy.

Steps

1. To enable FLR and searching backups, select **Enable indexing for file search and restore**.
2. To enable the debug logs for troubleshooting purposes, select **Troubleshooting**.
3. To override the default debug level, add the statement `debugLevel=<N>` to the `addon.cfg` configuration file, where `N` is the desired debug level, in the range `[4..9]`.

NOTE: Changing the default bug level can result in larger logs that can slow backup operations and sometimes cause the asset host to run out of disk space. On Windows hosts, the log files are located in the `C:\Program Files\DPSAPPS\fsagent\tmp` and `C:\Program Files\DPSAPPS\fsagent\logs` directories. On Linux hosts, the logs are located in the `/opt/dpsapps/fsagent/tmp` and `/opt/dpsapps/fsagent/logs` directories. If no backup jobs are running, files in these directories can be removed.

In Windows environments, the impacted logs include:

- `FSAgentInstallPath\logs\vescr.log`
- `FSAgentInstallPath\logs\nsriscsi.log`
- `FSAgentInstallPath\logs\nsriscsi_***.log`
- `FSAgentInstallPath\logs\nsrwriter.log`
- `FSAgentInstallPath\logs\ddfscn.***.log`
- `FSAgentInstallPath\logs\ddfscn_***.log`
- `FSAgentInstallPath\logs\ddfssv.log`
- `FSAgentInstallPath\logs\ddfssv_***.log`
- `FSAgentInstallPath\logs\ddfsrc_***.log`

In Linux environments, the impacted logs include:

- `/opt/dpsapps/fsagent/logs/nsriscsi.log`
- `/opt/dpsapps/fsagent/logs/ddfscn.***.log`
- `/opt/dpsapps/fsagent/logs/ddfssv.log`

If you have updated from an earlier File System agent version, some log files may appear with both `.log` and `.raw` extensions. Use the `.log` files.

4. Click **Next**.

Results

The wizard moves to the **Summary** page. Continue to Review the protection policy summary.

Review the protection policy summary

Review the protection policy group configuration details. You can click **Edit** next to any completed window's details to change any information. When completed, click **Finish**.

An informational message appears to confirm that PowerProtect Data Manager has saved the protection policy. For centralized backups, when a new protection policy is created, PowerProtect Data Manager performs the first full backup and subsequent backups according to the specified schedule.

Click **OK** to exit the window, or click **Go to Jobs** to open the **Jobs** window to monitor the backup of the new protection policy group.

Viewing protection policy jobs

You can monitor and view detailed information in the **Jobs** window for protection policy objectives, including backups and restores.

The **Cancel** and **Retry** options are not available for self-service jobs that are created by database application agents.

From the PowerProtect Data Manager UI left navigation pane, you can select **Jobs > Protection Jobs** to view the **Protection Jobs** window, which displays the protection job group status. You can also click the job ID in the **Protection Jobs** window to view the **Job ID Summary** window, which displays the status of each asset job.

The status of an asset job is **Skipped** when the asset is present in the asset host but unavailable for backup because it is offline or in a restoring, recovery pending, or suspect state. You can see the reason for the **Skipped** status in the details section of the **Job ID Summary** window.

When all the assets in a job group are skipped, the job group status appears as **Skipped** in the **Protection Jobs** window. When some but not all assets in a job group are skipped, the job group status appears as **Completed with Exceptions**.

When at least one asset in a job group has the **Failed** status, the job group status appears as **Failed**.

When a backup fails or a backup is skipped, the backup job steps appear as canceled for the particular database. The backup job steps are displayed on the **Step Log** tab in the details section of the **Job ID Summary** window.

Protection rules

Protection rules comprise one or more conditions that select matching assets and automatically assign them to a corresponding protection policy. PowerProtect Data Manager applies these rules to assets at discovery time.

You can apply protection rules to policies for the following asset types:

- Virtual machine
- File System
- Kubernetes
- Microsoft Exchange
- Microsoft SQL
- Network Attached Storage (NAS)
- Oracle
- SAP HANA
- PowerStore block volumes

Before defining a protection rule, note the following:

- Creating protection rules requires at least one existing protection policy.
- An asset can only belong to one protection policy.
- Assets can move from one policy to another policy based on the priorities of the protection rules. You can manually move an asset into a protection policy and override automatic placement through protection rules. Manual assignment protects the asset through the specified policy but protection rules no longer apply to that asset. To apply protection rules again, remove the asset from the protection policy.
- To ensure the protection of homogeneous assets, the protection rule must specify a storage asset type.
- For virtual machine protection policies, virtual machine tags created in the **vSphere Client** can only be applied to a protection rule.
- A virtual machine application-aware protection policy that protects a Microsoft SQL Server Always On availability group (AAG) must include all the virtual machines of the AAG in the same protection group. Failure to meet this requirement might result in Microsoft SQL Server transaction log backups being skipped. Ensure that the protection rules are designed to include all the AAG virtual machines.
- For Oracle assets, ensure that the Oracle protection rules do not use the DB ID and Oracle SID Name field settings that were supported with versions prior to PowerProtect Data Manager 19.6.

Protection rule attributes and criteria

The following table provides a list of the available rule attributes and criteria for all asset types that support the application of protection rules to policies. The available matching criteria depend on the selected attribute.

Table 11. Supported attributes and matching criteria by asset type

Asset type	Supported attributes	Matching criteria	Notes
Virtual machines	Cluster Name	Begins with, Contains, Does not contain, Does not equal, Ends with, Equals, Matches RegEx, Does Not Match RegEx	If using the Host Name for a virtual machine protection rule to determine which assets get included, ensure that you do not specify a host in a cluster. If you specify a host in a cluster, PowerProtect Data Manager will not protect the virtual machine assets under this host because although these assets are currently running within this host, they are not owned by the host and can be switched to another host under the same cluster at any time.
	Datacenter Name		
	Datastore Name		
	Host Name		
	OS Type		
	Power State	Equals, Does not equal	The Power State attribute is applicable only to virtual machines, and enables filtering of virtual machine hosts based on the state of the host.
	vCenter Name	Begins with, Contains, Does not contain, Does not equal, Ends with, Equals, Matches RegEx, Does Not Match RegEx	<ul style="list-style-type: none"> The VM Folder Name and VM Resource Pool attributes support protection for all virtual machine assets and resource pools in the selected folder and its subfolders. Regular expressions for the VM Folder Name and VM Resource Pool attributes use Google RE2J syntax. The operators and effects on the Optional tab of the dialog box are unavailable for these attributes. However, the operators and effects on the Unsupported tab are available, as are the standard regular expression predefined character classes. For example, \d for a digit. <p>Regular expressions for all other attributes use ElasticSearch regex syntax. These expressions do not support predefined character classes.</p> <p>Because predefined character classes are valid for some attributes, the UI does not mark these classes as invalid syntax. This is true even for attributes where such classes are not supported.</p>
	VM Display Name		
	VM Folder Name		
	VM Resource Pool		
VM Size	Greater than, Less than		
VM Tags	Equals, Does not equal, In, Not in		
File systems	File System Name	Begins with, Contains, Does not contain, Does not equal, Ends with, Equals, Matches RegEx, Does Not Match RegEx	
	Host Name		
	Cluster Name		
	Host Type		

Table 11. Supported attributes and matching criteria by asset type (continued)

Asset type	Supported attributes	Matching criteria	Notes
	File System Type	Begins with, Contains, Does not contain, Does not equal, Ends with, Equals	
	File System Size	Greater than, Less than	
Exchange servers	Application Name	Begins with, Contains, Does not contain, Does not equal, Ends with, Equals, Matches RegEx, Does Not Match RegEx	
	DB Name		
SQL databases	Application Name	Begins with, Contains, Does not contain, Does not equal, Ends with, Equals, Matches RegEx, Does Not Match RegEx	
	DB Name		
	Host/Cluster/Group Name		
	Host Type	Equals	
Oracle databases	DB Name	Begins with, Contains, Does not contain, Does not equal, Ends with, Equals, Matches RegEx, Does Not Match RegEx	The addition of a protection rule is not supported for a newly discovered Oracle Data Guard asset. When a PowerProtect Data Manager system is updated from an earlier version, the Oracle Data Guard asset continues to be associated with previously defined protection rules.
	Host/Cluster/Group Name		
	Host Type		
	Oracle Version	Begins with, Contains, Does not contain, Does not equal, Ends with, Equals, Matches RegEx, Does Not Match RegEx	
Kubernetes namespaces	Namespace Name	Begins with, Contains, Does not contain, Does not equal, Ends with, Equals, Matches RegEx, Does Not Match RegEx	
	Namespace Label	Equals, Does not equal	
SAP HANA databases	Application Name	Begins with, Contains, Does not contain, Does not equal, Ends with, Equals, Matches RegEx, Does Not Match RegEx	
	DB Name		
Network-attached storage	Share Name	Begins with, Contains, Does not contain, Does not equal, Ends with, Equals, Matches RegEx, Does Not Match RegEx	
	Share Size	Greater than, Less than	
	File System Name	Begins with, Contains, Does not contain, Does not equal, Ends with, Equals, Matches RegEx, Does Not Match RegEx	
	Protocol	Equals, Does not equal	

Table 11. Supported attributes and matching criteria by asset type (continued)

Asset type	Supported attributes	Matching criteria	Notes
	Server Name	Begins with, Contains,	
	Server Address	Does not contain, Does not equal, Ends with, Equals, Matches RegEx, Does Not Match RegEx	
Block Volumes	Block Volume Name	Begins with, Contains, Does not contain, Does not equal, Ends with, Equals, Matches RegEx, Does Not Match RegEx	
	Block Volume Size	Greater than, Less than	
	Block Volume WWN	Begins with, Contains, Does not contain, Does not equal, Ends with, Equals, Matches RegEx, Does Not Match RegEx	
	Origin Type	Equals, Does not equal	
	Replica Storage Array Name	Begins with, Contains, Does not contain, Does not equal, Ends with, Equals, Matches RegEx, Does Not Match RegEx	
	Storage Array IP	Begins with, Contains, Does not contain, Does not equal, Ends with, Equals, Matches RegEx, Does Not Match RegEx	
	Storage Array Name	Begins with, Contains, Does not contain, Does not equal, Ends with, Equals	

Add a protection rule

Select a protection policy and then define one or more conditions. Where applicable, create compound rules by linking multiple conditions through logical operators.

About this task

Compound rules enable you to combine multiple selection criteria through **AND** and **OR** operators for higher precision. For example, assets in a particular data center with particular tags. Compound rules must have at least one condition.

The **Add Protection Rule** wizard displays compound rules in containers. Grouping rules in the same container represents a logical AND of those rules. Placing rules in separate containers represent a logical OR of those rules. For example, the compound rule (A AND B) OR (C) corresponds to one container with rules A and B, and another container with rule C.

The wizard validates fields as you type. As you define the protection rule, the wizard also displays a count of assets which match the entire protection rule, next to **View Filtered Assets**.

Steps

1. From the PowerProtect Data Manager UI, select **Protection > Protection Rules**. The **Protection Rules** window appears.

2. Select the tab that corresponds to the asset type or host for which you would like to add the protection rule, and then click **Add**.

PowerProtect Data Manager supports the application of protection rules for the following asset types:

- Virtual machine
- File System (Block Volumes)
- Kubernetes namespaces
- Microsoft Exchange
- Microsoft SQL
- Network Attached Storage (NAS)
- Oracle
- SAP HANA
- PowerStore block volumes

The **Add Protection Rule** window opens to the **Select Protection Policy** page.

3. Select the target protection policy for the protection rule and then click **Next**.

The **Add Rule Description** page appears.


4. Define the purpose of the protection rule:

- a. **Name.** For example, **Rules Prod Finance**. The name must be unique.
- b. **Description.** For example, **Finance department production servers**
- c. Click **Next**.

The **Add Conditions** page appears:

5. Define the protection rule:

- a. Select an attribute. The available attributes depend on the selected host type.
- b. Select a matching criteria. The available matching criteria depend on the selected attribute.

i **NOTE:** Where the available matching criteria includes regular expressions, click  for a list of supported operators and effects in a separate dialog box.


- c. Depending on the selected attribute, supply a search phrase to compare against the attribute or select an option from the list.


The wizard displays a count of matching assets beside the rule and enables new **Add Rule** options for compound rules.

For example, a rule with the filters **VM Folder Name**, **Contains**, and **Finance** can match assets belonging to your finance department to the selected protection policy.


6. To define a compound rule:

The wizard only enables some **Add Rule** options after the successful validation of other rules in the same container. For example, rules cannot be empty.

- a. Select a logical operation, and then click the corresponding **Add Rule** option.
 - If you select **+ (AND)**, the new rule appears in the same container.
 - If you select **Add Rule - OR**, the new rule appears in a separate container.
- b. Repeat the previous step to define the new protection rule.
- c. To remove a rule from a compound rule, click  for that rule.

i **NOTE:** The wizard disables  for any rules whose deletion would result in an empty container. To remove these rules, remove the entire container.

The wizard removes the selected rule and any associated **Add Rule** options.

- d. To remove an entire container and any rules within it, click  for that container. The wizard also removes any associated **Add Rule** options.

- e. To remove all rules, click  **Reset Rules**.

The wizard displays a count of matching assets beside each rule and, for each container, a count of matching assets for all rules in the container.

i **NOTE:** The counts displayed by the **Protection > Protection Rules > Add Protection Rules > Add Conditions** and **Protection > Protection Rules > Add Protection Rules > Add Conditions > Filtered Assets** panes only count the number of assets in the filtered folders and resource pools. The counts do not include assets in subfolders or sub-resource pools. Despite the displayed count, all assets in subfolders and sub-resource pools are also protected. For existing protection rules, accurate asset counts are displayed in the **Protection > Protection Rules** and **Protection > Protection Policies** panes.

7. To see a list of unprotected assets which match the protection rule, click **View Matching Assets**. The **Matching Assets** window opens and displays the details of each matching asset. Verify that the list includes all expected assets, and then click **Done**.
8. If the protection rule and list of matching assets do not meet expectations, adjust the rules accordingly. Alternatively, reset the rules and then build the protection rule again.
9. If the protection rule and list of matching assets meet expectations, click **Next**. The **Summary** page appears.
10. Review the protection rule details and then click **Finish**.

Results

The new protection rule automatically protects any matching assets.

Manually run a protection rule

PowerProtect Data Manager automatically runs protection rules when new assets are detected or when existing assets are modified. You can also run protection rules manually.

Prerequisites

NOTE: For SQL, Oracle, SAP HANA, file system, and block volume asset types, the protection rule runs only on scheduled discovery in PowerProtect Data Manager. Ensure that you schedule discovery for these asset types.

Steps

1. From the PowerProtect Data Manager UI, select **Protection > Protection Rules**. The **Protection Rules** window appears.
2. Select the required protection rules, and then click **Run**. PowerProtect Data Manager runs all of the selected protection rules for the current asset type.

Schedule an asset source discovery

For some asset types, the protection rule runs only upon scheduled asset source discovery in PowerProtect Data Manager.

About this task

This task applies only to the following asset types:

- File System
- Microsoft SQL Server
- Block volumes
- SAP HANA

Steps

1. Select **Infrastructure > Asset Sources**.
2. Select the tab for the type of asset source that you want to discover.
3. Select the asset source name in the asset source list, and then click **Discover**.
4. From the **Discovery Schedule** list, select the time of day to initiate the discovery.

Edit or delete a protection rule

You can change the name, description, the rule filters, and the associated protection policy.

Steps

1. Select **Protection > Protection Rules**. The **Protection Rules** window appears.

NOTE: On the **Protection Rules** window, the **Source** column lists **Protection Policy** and **User** created rules. The **Protection Policy** rules, that is, the protection rules that are created dynamically while adding or editing the protection policy cannot be edited or deleted. Unselect the server/host of an asset by editing the protection policy to delete such rules.

2. To edit a protection rule, select the rule and then click **Edit**.

The **Edit Protection Rule** window appears.

- a. Select a protection policy, and then click **Next**.
 - b. Modify the name, description, or filter rules, and then click **Next**.
Add a protection rule provides more information about working with rules.
 - c. Review the protection rule summary, and then click **Finish**.
3. To delete a protection rule, select the rule and then click **Delete**.

PowerProtect Data Manager removes from protection policies any assets that were added because of this protection rule. PowerProtect Data Manager adds those assets again if you do not update related protection rules.

View assets applied to a protection rule

You can view the assets that are applied to a protection rule from the **Protection Rules** window. If the modification of a protection rule results in assets moving from one policy to another, the **Protection Rules** window enables you to verify the results.

About this task

To view assets that are applied to a protection rule, complete the following steps.

Steps

1. From the left navigation pane, select **Protection > Protection Rules**.
The **Protection Rules** window appears.
2. Click the link in the **Assigned Assets Count** column for the protection rule.
The **Assets List** window appears and displays the matched assets.
3. To export asset records for the protection rule, in the **Assets List** window, click **Export All**.

Change the priority of an existing protection rule

When multiple protection rules exist, you can define the priority of each rule. Priority determines which rule applies to an asset when that asset matches multiple rules and those rules have conflicting actions.

About this task

For example, if an asset matches several protection rules and each rule specifies a different protection policy, then the rule with the highest priority determines the policy assignment.

Protection rule priorities are integers. Smaller integers represent a higher priority.

Steps

1. Select **Protection > Protection Rules**.
The **Protection Rules** window appears.
2. To change a protection rule's priority, select the rule and then click **Up** or **Down**.
Remember that the smaller integer has the higher priority.

Configure protection rule behavior

You can use the REST API to configure what happens when a protection rule changes.

The PowerProtect Data Manager Public REST API documentation provides instructions:

NOTE:

If you update from a previous release of PowerProtect Data Manager, the configured behavior for protection rule changes still applies to the current release. For example, in PowerProtect Data Manager 19.12, if you did not configure protection rules through `application.properties` to move assets across policies, then you cannot change the behavior with this method in PowerProtect Data Manager 19.13 or later.

However, if you updated the configuration file to enable protection rules to move assets across policies, then this behavior continues to apply after the update.

Cancel a File System agent backup or restore job

You can cancel a File System agent protection job (backup or restore) from the PowerProtect Data Manager UI. The job must be in a queued or running state. The backup or restore job runs for a primary backup that is configured through a File System agent protection policy.

About this task

You can perform two types of application agent job cancellations in the PowerProtect Data Manager UI:

- Cancellation of a job group that includes one or more asset jobs.
- Cancellation of an individual asset job.

NOTE:

- On a Linux platform, if a block-based image restore fails, or if you cancel a block-based image restore while it is Running, you must manually mount the target volume before next attempting any backup or restore on the same volume.
- Upon cancellation of an incremental block-based backup, the next backup is promoted automatically to a full backup.
- When a job completes before the cancel request reaches the application host, the status of the canceled job transitions to either success or failure.
- You can cancel many other types of jobs, in addition to protection jobs. The *PowerProtect Data Manager Administration and User Guide* provides more information.

Perform the following steps to cancel an application agent protection job in the PowerProtect Data Manager UI.

Steps

1. From the PowerProtect Data Manager UI left navigation pane, select **Jobs > Protection Jobs**.

The **Protection Jobs** window opens to display a list of protection jobs and job groups.

2. In the **Protection Jobs** window, perform the required type of job cancellation:

- To cancel a job group:
 - a. In the **Protection Jobs** window, select the required job group and click **Cancel**.
A job group warning prompt appears.
 - b. Click **OK** at the prompt.

You can monitor the job group cancellation in the **Protection Jobs** window. The job group status changes to Canceled when the cancellation of all the asset jobs is complete.

To monitor the cancellation of individual asset jobs within the job group, click the job ID in the **Protection Jobs** window. The **Job ID Summary** window opens, where you can view the status of each asset job.

- To cancel an asset job:
 - a. In the **Protection Jobs** window, click the job ID.
The **Job ID Summary** window opens to display the job details of the assets in the job group.
 - b. In the **Job ID Summary** window, select the required asset job and click **Cancel**.
A job warning prompt appears.
 - c. Click **OK** at the prompt.

You can monitor the asset job cancellation in the **Job ID Summary** window. The asset job status changes to **Cancelled** when the job cancellation is complete.

NOTE: When the cancel request for a job cannot be completed, an informational alert is displayed.

Add a service-level agreement

SLA Compliance in the PowerProtect Data Manager UI enables you to add a service-level agreement (SLA) that identifies your service-level objectives (SLOs). You use the SLOs to verify that your protected assets are meeting the service-level agreements (SLAs).

About this task

NOTE: When you create an SLA for Cloud Tier, you can include only full backups in the SLA.

NOTE: The **Extended Retention** SLA only applies to protection policies created in PowerProtect Data Manager 19.11 or earlier. The Extended Retention objective was removed in PowerProtect Data Manager 19.12. Protection policies that were created in earlier releases with the **Extended Retention** SLA are supported. However, you cannot edit the **Extended Retention** SLA in these policies.

In the **SLA Compliance** window, you can export compliance data by using the **Export All** functionality.

Steps

1. From the PowerProtect Data Manager UI, select **Protection > SLA Compliance**.

The **SLA Compliance** window appears.

2. Click **Add** or, if the assets that you want to apply the SLA to are listed, select these assets and then click **Add**.

The **Add Service Level Agreement** wizard appears.

3. Select the type of SLA that you want to add, and then click **Next**.

- **Policy.** If you choose this type, go to step 4.
- **Backup.** If you choose this type, go to step 5.
- **Replication.** If you choose this type, go to step 6.
- **Cloud Tier.** If you choose this type, go to step 7.

You can select only one type of Service Level Agreement.

4. If you selected **Policy**, specify the following fields regarding the purpose of the new Policy SLA:

- a. The **SLA Name**.
- b. If applicable, select **Minimum Copies**, and specify the number of Backup, Replication, and Cloud Tier copies.
- c. If applicable, select **Maximum Copies**, and specify the number of Backup, Replication, and Cloud Tier copies.
- d. If applicable, select **Available Location** and select the applicable locations. To add a location, click **Add Location**.

Options include the following:

- **In**—Include locations of all copies in the SLO locations. Selecting this option does not require every SLO location to have a copy.
- **Must In**—Include locations of all copies in the SLO locations. Selecting this option requires every SLO location to have at least one copy.
- **Exclude**—Locations of all copies must be non-SLO locations.

NOTE: Policy files backed up on a storage unit with indefinite retention hold (IRH) enabled cannot be deleted or modified, even after retention lock expiry. It is therefore recommended that you do not select the **Maximum Copies** option because this setting conflicts with IRH. Otherwise, the SLA will not complete successfully once the number of copies exceeds the specified number.

- e. If applicable, select **Allowed in Cloud through Cloud Tier/Cloud DR**.
- f. Click **Finish**, and then go to step 9.

5. If you selected **Backup**, specify the following fields regarding the purpose of the new **Backup** SLA:

- a. The **SLA Name**.

- b. If applicable, select **Recovery Point Objective required (RPO)**, and then set the duration. The purpose of an RPO is business continuity planning. It indicates the maximum targeted period in which data (transactions) might be lost from an IT service due to a major incident.
- NOTE:** You can select only **Recovery Point Objective required** to configure as an independent objective in the SLA, or select both **Recovery Point Objective required** and **Compliance Window for copy type**. If you select both, the RPO setting must be one of the following:
- Greater than 24 hours or more than the Compliance window duration, in which case RPO validation occurs independent of the Compliance Window.
 - Less than or equal to the Compliance Window duration, in which case RPO validation occurs within the Compliance Window.
- c. If applicable, select **Compliance Window for copy type**, and then select a schedule level from the list, for example, **All, Full, Cumulative**, and set the duration. **Duration** indicates the amount of time necessary to create the backup copy. Ensure that the **Start Time** and **End Time** of backup copy creation falls within the Compliance Window duration specified.
- This window specifies the time during which you expect the specified activity to take place. Any specified activity that occurs outside of this **Start Time** and **End Time** triggers an alert.
- d. If applicable, select the **Verify expired copies are deleted** option.
- Verify expired copies are deleted** is a compliance check to see if PowerProtect Data Manager is deleting expired copies. This option is disabled by default.
- NOTE:** Data that is backed up on a storage unit with IRH enabled cannot be deleted or modified, even after retention lock expiry. It is therefore recommended that you do not select the **Verify expired copies are deleted** option because this setting conflicts with IRH. Otherwise, the SLA will not complete successfully.
- e. If applicable, select **Retention Time Objective**, and specify the number of Days, Months, Weeks, or Years.
- NOTE:** The value of **Retention Time Objective** must match the lowest retention value of the backup levels of the target objectives of this policy. For example, if the synthetic full backup **Retain For** is 30 days but the full backup **Retain For** is 60 days, set the **Retention Time Objective** to 30 days.
- f. If applicable, select the **Verify Retention Lock is enabled for all copies** option. This option is disabled by default.
- g. Click **Finish**, and go to step 9.
The **SLA Compliance** window appears with the new SLA.
6. If you selected **Replication**, specify the following fields regarding the purpose of the new Replication SLA:
- a. The **SLA Name**.
- b. If applicable, select the **Compliance Window**, and specify the **Start Time** and **End Time**.
- This window specifies the times that are permissible and during which you can expect the specified activity to occur. Any specified activity that occurs outside of this start time and end time triggers an alert.
- c. If applicable, select the **Verify expired copies are deleted** option.
- Verify expired copies are deleted** is a compliance check to see if PowerProtect Data Manager is deleting expired copies. This option is disabled by default.
- NOTE:** Data that is replicated on a storage unit with IRH enabled cannot be deleted or modified, even after retention lock expiry. It is therefore recommended that you do not select the **Verify expired copies are deleted** option because this setting conflicts with IRH. Otherwise, the SLA will not complete successfully.
- d. If applicable, select **Retention Time Objective**, and specify the number of Days, Months, Weeks, or Years.
- NOTE:** Set the value of **Retention Time Objective** to match the lowest retention value of the backup levels of the target objectives of this policy.
- e. If applicable, select the **Verify Retention Lock is enabled for all copies** option. This option is disabled by default.
- f. Click **Finish**, and go to step 9.
The **SLA Compliance** window appears with the newly added SLA.
7. If you selected **Cloud Tier** type SLA, specify the following fields regarding the purpose of the new Cloud Tier SLA:
- a. The **SLA Name**.
- b. If applicable, select the **Verify expired copies are deleted** option.
- This option is a compliance check to determine if PowerProtect Data Manager is deleting expired copies. This option is disabled by default.
- c. If applicable, select **Retention Time Objective** and specify the number of Days, Months, Weeks, or Years.

NOTE: Set the value of **Retention Time Objective** to match the lowest retention value of the backup levels of the target objectives of this policy.

- d. If applicable, select the **Verify Retention Lock is enabled for all copies** option. This option is disabled by default.
 - e. Click **Finish**.
8. If the SLA has not already been applied to a protection policy:
- a. Go to **Protection > Protection Policies**.
 - b. Select the policy, and then click **Edit**.
9. In the **Objectives** row of the **Summary** window, click **Edit**.
10. Do one of the following, and then click **Next**:
- Select the added Policy SLA from the **Set Policy Level SLA** list.
 - Create and add the SLA policy from the **Set Policy Level SLA** list.
- The **Summary** window appears.

11. Click **Finish**.

An informational message appears to confirm that PowerProtect Data Manager has saved the protection policy.

12. Click **Go to Jobs** to open the **Jobs** window to monitor the backup and compliance results, or click **OK** to exit.

NOTE: Compliance checks occur automatically every day at 2 a.m. Coordinated Universal Time (UTC). If any objectives are out of compliance, an alert is generated at 2 a.m. UTC. The **Validate** job in the **System Jobs** window indicates the results of the daily compliance check.

For a backup SLA with a required RPO setting that is less than 24 hours, PowerProtect Data Manager performs real-time compliance checks. If you selected **Compliance Window for copy type** and set the backup level to **All**, the real-time compliance check occurs every 15 minutes only within the compliance window. If the backup level is not **All**, or if a compliance window is not specified, the real-time compliance check occurs every 15 minutes without stop.

NOTE: If the backup SLA has a required RPO setting of 24 hours or greater, compliance checks occur daily at 2 a.m. UTC. Real-time compliance checks do not occur for backup SLAs with an RPO setting of 24 hours or greater.

Real-time compliance-check behavior

If the asset was not backed up within the RPO backup interval requirement, an alert indicates that the RPO of the asset is out of compliance. This alert is generated once within an RPO period. If the same backup copy is missed when the next compliance check occurs, no further alerts are generated.

If the asset was backed up within the RPO backup interval requirement, the RPO of the asset is in compliance.

If multiple assets in a policy are out of compliance, a single alert is generated. This alert includes information for all the assets in the policy. In the **Alerts** window, the asset count next to the alert summary indicates the number of assets that are out of compliance in the policy.

13. In the **Jobs** window, click  next to an entry to view details on the SLA Compliance result.

Extended retention for protection policies created in PowerProtect Data Manager 19.11 or earlier

NOTE: This section only applies to protection policies created in PowerProtect Data Manager 19.11 or earlier. For protection policies created in PowerProtect Data Manager 19.12 or later, you add multiple full schedules for primary backup and replication objectives. Protection policies that were created in an earlier release with the **Extend Retention** objective are supported. However, you cannot edit existing extended retention objectives or add new extended retention objectives in these policies. Knowledge Base article 000204454 at <https://www.dell.com/support/> provides detailed information about specific **Extend Retention** objective migration scenarios when updating PowerProtect Data Manager.

For protection policies created in PowerProtect Data Manager 19.11 or earlier, the **Extend Retention** objective allows you to extend the retention period for the primary backup copy for long-term retention. For example, your regular schedule for daily backups uses a retention period of 30 days. However, you can extend the retention period to keep full backups taken on Mondays for 10 weeks.

Both centralized and self-service protection policies support weekly, monthly, and yearly recurrence schedules to meet the demands of your compliance objectives. For example, you can retain the last full backup containing the last transaction of a

fiscal year for 10 years. Extended retention periods can retain scheduled full backups with a repeating pattern for a specified amount of time.

For example:

- Retain full yearly backups that are set to repeat on the first day of January for 5 years.
- Retain full monthly backups that are set to repeat on the last day of every month for 1 year.
- Retain full yearly backups that are set to repeat on the third Monday of December for 7 years.

Preferred alternatives

When you define an extended retention objective for a protection policy, you define matching criteria that select preferred backups to retain. If the matching criteria do not identify a matching backup, PowerProtect Data Manager automatically retains the preferred alternative backup according to one of the following methods:

- Look-back—Retain the last available full backup that was taken before the matching criteria.
- Look-forward—Retain the next available full backup that was taken after the matching criteria.

For example, consider a situation where you configured a protection policy to retain the daily backup for the last day of the month to extended retention. However, a network issue caused that backup to fail. In this case, look-back matching retains the backup that was taken the previous day, while look-forward matching retains the backup that was taken the following day.

By default, PowerProtect Data Manager uses look-back matching to select the preferred alternative backup. A grace period defines how far PowerProtect Data Manager can look in the configured direction for an alternative backup. If PowerProtect Data Manager cannot find an alternative backup within the grace period, extended retention fails.

You can use the REST API to change the matching method or the grace period for look-forward matching. The PowerProtect Data Manager Public REST API documentation provides instructions. If there are no available backups for the defined matching period, you can change the matching method to a different backup.

For look-forward matching, the next available backup can be a manual backup or the next scheduled backup.

Selecting backups by weekday

This section applies to centralized protection policies. Self-service protection policies have no primary backup objective configuration.

When you configure extended retention to match backups by weekday, PowerProtect Data Manager might identify a backup as having been taken on the wrong weekday. This behavior happens where the backup window does not align with the start of the day. PowerProtect Data Manager identifies backups according to the day on which the corresponding backup window started, rather than the start of the backup itself.

For example, consider a backup schedule with an 8:00 p.m. to 6:00 a.m. backup window:

- Backups that start at 12:00 a.m. on Sunday and end at 6:00 a.m. on Sunday are identified as Saturday backups, since the backup window started on Saturday.
- Backups that start at 8:01 p.m. on Sunday and end at 12:00 a.m. on Monday are identified as Sunday backups, since the backup window started on Sunday.
- Backups that start at 12:00 a.m. on Monday and end at 6:00 a.m. on Monday are identified as Sunday backups, since the backup window started on Sunday.

In this example, when you select Sunday backups for extended retention, PowerProtect Data Manager does not retain backups that were taken between 12:00 a.m. and 8:00 p.m. This behavior happens even though the backups occurred on Sunday. Instead, PowerProtect Data Manager selects the first available backup that started after 8:00 p.m. on Sunday for extended retention.

If no backups were created between 8:01 p.m. on Sunday and 6:00 a.m. on Monday, PowerProtect Data Manager retains the next alternative to extended retention. In this example, the alternative was taken after 8:00 a.m. on Monday.

Extended retention backup behavior

When PowerProtect Data Manager identifies a matching backup, automatically extended retention creates a job at the beginning of the backup window for the primary objective. This job remains queued until the end of the backup window.

The following examples describe the behavior of backups with extended retention for centralized and self-service protection.

Centralized protection

An hourly primary backup schedule starts on Sunday at 6:00 p.m., and ends on Monday at 6:00 p.m. with a weekly extended retention objective set to repeat every Sunday. PowerProtect Data Manager selects the first available backup starting after 8:00 p.m. on Sunday for long-term retention.

The following diagram illustrates the behavior of backups with extended retention for a configured protection policy. In this example, full daily backups starting at 10:00 p.m. and ending at 6:00 a.m. are kept for 1 week. Full weekly backups are set to repeat every Sunday and are kept for 1 month.




Figure 1. Extend retention backup behavior

Self-service protection

For self-service backups, PowerProtect Data Manager uses a default backup window of 24 hours. A backup schedule starts on Sunday at 12:00 p.m. and ends on Monday at 12:00 p.m. with a weekly extended retention objective set to repeat every Sunday. PowerProtect Data Manager selects the first available backup that is taken between 12:00 p.m. on Sunday and 12:00 p.m. on Monday for long-term retention.

Replication of extended retention backups

You can change the retention time of selected full primary backups in a replication objective by adding a replication objective to the extended retention backup. The rules in the extended retention objective define the selected full primary backups. Review the following information about replication of extended retention backups.

- Before you configure replication of extended retention backups, create a replication objective for the primary backup.
- Configure the replication objective of the extended retention and match this objective with one of the existing replication objectives based on the primary backup. Any changes to a new or existing storage unit in the extended retention replication objective or the replication objective of the primary backup is applied to both replication objectives.
- The replication objective of extended retention backups only updates the retention time of replicated backup copies. New backup copies are not created in the replication storage.

Edit the retention period for backup copies



You can edit the retention period of one or more backup copies to extend or shorten the amount of time that backups are retained.

About this task

You can edit the retention period for all asset types and backup types.

Steps

1. Select **Infrastructure > Assets**.

- From the **Assets** window, select the tab for the asset type for which you want to edit the retention period. If a policy has been assigned, the table lists the assets that have been discovered, along with the associated protection policy.
 - Select a protected asset from the table, and then click **View Copies**. The **Copy Locations** pane identifies where the backups are stored.
 - In the left pane, click  to the right of the icon for the asset. The table in the right pane lists the backup copies.
 - Select one or more backup copies from the table, and click **Edit Retention**.
 - Select one of the following options:
 - To select a calendar date as the expiration date for backups, select **Retention Date**.
 - To define a fixed retention period in days, weeks, months, or years after the backup is performed, select **Retention Value**. For example, you can specify that backups expire after 6 months.
-  **NOTE:** When you edit the retention period for copies that are retention locked, you can only extend the retention period.
- When satisfied with the changes, click **Save**.
The asset is displayed in the list with the changes. The **Retention** column displays both the original and new retention periods, and indicates whether the retention period has been extended or shortened.

Delete backup copies


In addition to deleting backups after the retention period expires, PowerProtect Data Manager enables you to manually delete backup copies from protection storages.

About this task

If you no longer require a backup copy and the retention lock is not enabled, you can delete backup copies prior to their expiration date.


You can perform a backup copy deletion that deletes only a specified part of a backup copy chain, without impacting the ability to restore other backup copies in the chain. When you select a specific backup copy for deletion, only that backup copy and the backup copies that depend on the selected backup copy are deleted. For example, when you select to delete a full backup copy, any other backup copies that depend on the full backup copy are also deleted.

Steps

- From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
- From the **Assets** window, select the tab for the asset type for which you want to delete copies. If a policy has been assigned, the table lists the assets that have been discovered, along with the associated protection policy.
- Select a protected asset from the table, and then click **View Copies**. The **Copy Locations** pane identifies where the backups are stored.
- In the left pane, click  to the right of the icon for the asset. The table in the right pane lists the backup copies.
- Select one or more copies from the table that you want to delete from the DD system, and then click **Delete**.
A preview window opens and displays the selected backup copies.
- For all asset types, you can choose to keep the latest backup copies or delete them. By default, PowerProtect Data Manager keeps the latest backup copies. To delete the latest backup copies, clear the check box next to **Include latest copies**.
- To delete the backup copies, in the preview window, click **Delete**.

 **NOTE:** The delete operation may take a few minutes and cannot be undone.

An informational dialog box opens to confirm the copies are being deleted. To monitor the progress of the operation, click **Go to Jobs**. To view the list of backup copies and their status, click **OK**.

 **NOTE:** If the data deletion is successful but the catalog deletion is unsuccessful, then the overall deletion job status appears as **Completed with Exceptions**.

When the job completes, the task summary provides details of each deleted backup copy, including the time that each copy was created, the backup level, and the retention time. The time of copy creation and the retention time are shown in UTC.


An audit log is also generated and provides details of each deleted backup copy, including the time that each copy was created, the backup level, and the retention time. The time of copy creation and the retention time are shown in UTC. Go to **Alerts > Audit Logs** to view the audit log.

- Verify that the copies are deleted successfully from protection storage. If the deletion is successful, the deleted copies no longer appear in the table.

Retry a failed backup copy deletion

If a backup copy is not deleted successfully, you can manually retry the operation.


Steps

- From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
- From the **Assets** window, select the tab for the asset type for which you want to delete copies. If a policy has been assigned, the table lists the assets that have been discovered, along with the associated protection policy.
- Select a protected asset from the table, and then click **View Copies**. The **Copy Locations** pane identifies where the backups are stored.
- In the left pane, click  to the right of the icon for the asset. The table in the right pane lists the backup copies.
- Select one or more backup copies with the **Deletion Failed** status from the table, and then click **Delete**.
You can also filter and sort the list of backup copies by status in the **Copy Status** column.
The system displays a warning to confirm that you want to delete the selected backup copies.
- Click **OK**.
An informational dialog box opens to confirm that the copies are being deleted. To monitor the progress of the operation, click **Go to Jobs**. To view the list of backup copies and their status, click **OK**.
- Verify that the copies are successfully deleted from protection storage. If the deletion is successful, the deleted copies no longer appear in the table.

Export data for deleted Exchange, File System, Kubernetes, Block Volume, and SQL backup copies

This option enables you to export results of deleted backup copies to a .csv file so that you can download an Excel file of the data.

Steps

- From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
- From the **Assets** window, select the tab for the asset type for which you want to export results of deleted backup copies. If a policy has been assigned, the table lists the assets that have been discovered, along with the associated protection policy.
- Select one or more protected assets from the table, and then select **More Actions > Export Deleted Copies**.
If you do not select an asset, PowerProtect Data Manager exports the data for deleted backup copies for all assets for the specific asset type.
- Specify the following fields for the export:
 - Time Range**
The default is **Last 24 Hours**.
 - Copy Status**
In order to export data for deleted backup copies, the backup copies must be in one of the following states:
 - Deleted**—The copy is deleted successfully from protection storage, and, if applicable, the agent catalog is deleted successfully from the agent host.
 - Deleting**—Copy deletion is in progress.
 - Deletion Failed**—Copy deletion from protection storage is unsuccessful.
-  **NOTE:** You cannot export data for backup copies that are in an **Available** state.
- Click **Download**.
If applicable, the navigation window appears for you to select the location to save the .csv file.
- Save the .csv file in the desired location and click **Save**.


Remove Exchange, File System, Kubernetes, Block Volume, and SQL backup copies from the PowerProtect Data Manager database

This option enables you to delete the backup copy records from the PowerProtect Data Manager database, but keep the backup copies in protection storage.

About this task

For backup copies that could not be deleted from protection storage, you can remove the backup copies from the PowerProtect Data Manager database. Removing the backup copies from PowerProtect Data Manager does not delete the copies in protection storage.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
2. From the **Assets** window, select the tab for the asset type for which you want to delete copies. If a policy has been assigned, the table lists the assets that have been discovered, along with the associated protection policy.
3. Select a protected asset from the table, and then click **View Copies**. The **Copy Locations** pane identifies where the backups are stored.
4. In the left pane, click  to the right of the icon for the asset. The table in the right pane lists the backup copies.
5. Select one or more backup copies with the **Deletion Failed** status from the table, and then click **Remove from PowerProtect**.
The system displays a warning to confirm that you want to delete the selected backup copies.
6. Click **OK**.
An informational dialog box opens to confirm that the copies are being deleted. To monitor the progress of the operation, click **Go to Jobs**. To view the list of backup copies and their status, click **OK**.
7. Verify that the copies are deleted from the PowerProtect Data Manager database. If the deletion is successful, the deleted copies no longer appear in the table. The backup copies remain in protection storage.

Host CPU throttling

A user with an administrator or backup role can enable the host CPU throttling feature to reduce the amount of CPU utilization by backups on the application agent host. The host CPU throttling can slow down the running backup jobs on the host and prevent any new backup jobs from starting. This feature is useful when backups on the host use a significant amount of CPU that adversely affects the existing production operations.

The user can enable the CPU throttling in the PowerProtect Data Manager UI and specify a host CPU utilization threshold or limit. When CPU utilization on the host exceeds the limit, the running backups are slowed down and new backups are queued. As a result, other processes on the same host can increase their CPU utilization and improve their response to critical tasks.

The host CPU throttling feature is supported for centralized and self-service backups with version 19.14 and later of the File System agent, Microsoft SQL Server agent, and Oracle RMAN agent. By default, CPU throttling is not enabled on an application agent host.

NOTE:

Use CPU throttling with care, on only specific hosts and in specific situations, as it can slow down the backup jobs by up to 20 times. The slowdown can cause backup jobs to run longer than the execution time window and possibly affect the compliance associated with the policies on the host.

The host CPU throttling feature is supported for Oracle Incremental Merge new backups. The feature is not supported for Oracle Incremental Merge existing in-progress centralized backups, self-service backups, or stand-alone application agents. You can enable CPU throttling for self-service Oracle RMAN agent backups through the following type of setting in the RMAN script. The backups are then slowed down:

```
ALLOCATE CHANNEL C1 TYPE SBT TRACE 5 PARMS
'SBT_LIBRARY=/home/oracle/opt/dpsapps/rmanagent/lib/libddobk.so,
SBT_PARMS=(RMAN_AGENT_HOME=/home/oracle/opt/dpsapps/rmanagent/,
STORAGE_UNIT=/F1C-RAC-DG-blrv136g138-e8e4e, BACKUP_HOST=10.125.208.242,
ORACLE_HOME=/u02/software, SCHEDULED_BACKUP=TRUE, CPU_THROTTLE=90)';
```

For example, a user with an administrator or backup role enables the host CPU throttling and specifies the CPU utilization limit of 90% on the application agent host. The effect on the backup jobs is as follows:

- When the host CPU utilization exceeds 90%:
 - Existing backup jobs are slowed down.
 - New backup jobs are queued on the host.
- When the host CPU utilization is reduced to 90% or less:
 - The slowed backup jobs resume their normal speed.
 - The queued backup jobs are started.

NOTE:

It is recommended that you keep the host CPU utilization limit above 85%. Setting a limit of 85% or less can significantly slow down the backup operations. If all backup jobs become queued, then disable the CPU throttling or change the CPU utilization limit as it was likely set too low.

Running backup jobs are not affected by a change to the host CPU utilization limit. A running backup continues to use the limit value from the time that the backup job started. Backup jobs can be canceled when throttled or queued.

If the host CPU usage is consistently high without running backups, the CPU throttling feature will not be effective and backups will likely never start after you set the CPU utilization limit. In this case, reduce the load on the host first and then use the host CPU throttling.

If the host includes pre-19.14 application agents or agents other than those for File System, Microsoft SQL Server, or Oracle RMAN, host CPU throttling is enabled, and the CPU utilization limit is exceeded:

- New backup jobs are queued on the host.
- Running backup jobs are not throttled.

A user can change the CPU utilization limit setting on a host as long as the host has a supported application agent. When CPU throttling is enabled on the host and then all the supported application agents are removed, the only available option for the user is to disable the CPU throttling.

Enable the host CPU throttling

To set the CPU throttling at the host level, perform the following steps as a user with an administrator or backup role.

1. From the left navigation pane in the PowerProtect Data Manager UI, select **Infrastructure > Application Agents**.

Starting with version 19.14, the lower table in the **Application Agents** window has two new columns that are related to host CPU throttling:

- The **Throttling Status** column displays the status of backup throttling on the application agent host:
 - **No Throttling**—CPU throttling is not set for backups on the host.
 - **Throttling**—CPU throttling is set for backups on the host.
 - **Unsupported**—CPU throttling is unsupported because the host has only pre-19.14 application agents or application agents other than the File System agent, Microsoft SQL Server agent, or Oracle RMAN agent.
- The **CPU Throttling** column displays the CPU utilization limit setting for backup throttling on the application agent host. The displayed setting is either a double dash when throttling is not set or the CPU utilization limit as a percentage between 1 to 99, inclusive.

By default, the **CPU Throttling** column is hidden. To display the column, click the icon on the lower left and then select **CPU Throttling** from the **Show/Hide Columns** list.

View application agent details provides more information about the **Application Agents** window.

2. To enable the backup CPU throttling for the host and set the CPU utilization limit:
 - a. In the lower table in the **Application Agents** window, select the application agent host.
 - b. Select **More Actions > Backup Throttling**.
 - c. In the **Set Backup Throttling** dialog box, specify the following fields and click **Save**. A system job is then created for the host CPU throttling configuration. When the CPU throttling limit is changed, an audit is also created:
 - **Backup Throttling**—Set the value to **Enabled**.
 - **CPU Utilization Limit**—Set the value to an integer between 1 and 99, inclusive. For example, set the value to 90 for a 90% limit. It is recommended that you keep the CPU utilization limit above 85%.

The specified CPU utilization limit appears as a percent value in the **CPU Throttling** column in the table in the **Application Agents** window.

When backup throttling occurs, a message appears in the step log that is available through the protection job "details" view. The message states that the backup speed has slowed due to application agent throttling restrictions. When a backup job is queued during CPU throttling, the following message appears in the step log:

```
Queued due to host backup throttling
```

On the **Details** tab in the protection job window, the value Yes appears after the Throttled field label when throttling has occurred at some point during the backup job.

Exclusion filters for File Systems

Exclusion filters enable you to exclude certain files and folders from protection, based on the filter's conditions (conditions for exclusion).

Use the PowerProtect Data Manager UI to add, edit, and delete exclusion filters for file system files and folders.

When you create or edit a protection policy, you can apply exclusion filters to the protection policy.

When an exclusion filter is applied to a protection policy, the File System agent performs file-based backups of the protected assets. File-based backups traverse through the entire directory structure of the file system to back up all the files in each directory of the file system. While file-based backups can provide additional capabilities such as exclusion, these backups take longer to complete when compared to block-based backups.

NOTE:

- Exclusion filters cannot be applied to self-service protection policies or to backups taken through self-service CLI.
- If an exclusion filter is added to a protection policy, file-based backups will be performed instead of block-based backups.

Add an exclusion filter

Use the PowerProtect Data Manager UI to add filters that exclude specific files and folders based on certain conditions, such as file type, file size, modification time, and file path. When a file or folder meets the conditions, the filter excludes the data from the backup for the protection policy.

About this task

Use this procedure to add up to four filters for a file.

Steps

1. Select **Protection > File Exclusion**.
The **File Exclusion** window appears.
2. Click **ADD**.
The **Filter Information** window appears.
3. In the **Filter Name** field, type a name for the filter.
4. In the **Description** field, describe the purpose of the filter.
5. Select a filtering condition. You can add multiple filters.

The filter excludes all files and folders that match these criteria from the backup for the protection policy. When you add multiple conditions, a file is excluded only if it meets all filter conditions. Within a filter, you can add a condition only once.

Available filtering conditions:

File Size	Exclude files and folders that are larger, smaller, or equal to a specified size. Specify a value in the Greater than or Equal to , Less than or Equal to , or Equal to field.
File Type	Exclude files or folders based on file type. Specify a file name extension or multiple file name extensions that are separated by commas.

Modified Time Exclude files or folders that were modified before or after a certain date. Specify a date in either the **After** or **Before** field.

NOTE:

- Ensure that the Network Time Protocol (NTP) server is properly configured to use the **Modified Time** filter.
- For NAS shares, enter the modified time in the Coordinated Universal Time (UTC) zone. If a different time zone is used, the system will still consider it in UTC, which might exclude incorrect folders/files.

Folder Path Exclude files and folders in a specific path. You can specify an absolute or relative path.

6. When you are finished building the filter, click **Add Filters**. The new filter appears in the table.
7. You can add up to four filters using the previous steps. When you are finished, click **Next**.
8. In the **Summary Confirmation** page, verify the filter information and click **Finish**.

Guidelines for exclusion filters for File Systems

Review the following guidelines for exclusion filters.

Excluding by file size

When viewing the size of a file on a remote file system, only the value in bytes is precise. The calculations used to round the bytes to a different measurement are specific to each vendor and operating system, and should not be used to exclude a specific file.

For example, a Windows host could indicate the following information for the size of a file:

Size: 750 MB (786,652,672 bytes)

If you use an exclusion filter of **Equal to** with a value of 750 MB, the rounding of bytes to megabytes will almost certainly result in the file not being excluded.

Examples of exclusion filters to exclude this file by size include the following:

- **Equal to** with a value of 786,652,672 bytes.
- **Greater than or Equal to** with a value of 749 MB or 785,000,000 bytes.
- **Less than or Equal to** with a value of 751 MB or 787,000,000 bytes.

Using wildcards

Supported wildcards include an asterisk (*) to represent zero or more characters and a question mark (?) to represent zero or one character.

NOTE: Be careful when using the wildcard *. Depending on the wildcard location, you can match folders whose name matches the filter pattern and their contents, even when the names of those files do not match the filter. For example, `*\log*.txt` also excludes files with the `.txt` extension in a folder whose name starts with `log`, even if the names of the files do not start with `log`.

Unsupported characters in file and path names

File and path names that contain a comma (,) cannot be directly specified in a filter. To exclude such a file or path, use a wildcard.

Excluding by file type

The **File Type** filter enables you to exclude files and folders based on file extension.

You can specify a single extension or multiple file extensions. Separate multiple entries with a comma and do not add a space between entries. You can also specify related extensions by using wildcards. For example, *.doc? matches both .doc files and .docx files.

Excluding by type and path

You can combine extension and path to exclude all files of a particular type without respect to the file location.

For example *log*.txt matches all text files (.txt) where the file name starts with log. at any path.

You can also exclude all files of a particular type from a specific path or multiple paths that are separated by commas. For example,

- Specifying C:\abc*.txt excludes all the text files in the folder C:\abc.
- Specifying C:\folder1*.txt, D:\folder2*.doc excludes all the text files in the C:\folder1 and the doc files in the D:\folder2 folders.

NOTE: All the matching files under subfolders of the specified path are recursively excluded.

You can combine these guidelines to exclude all files that match a specific name pattern under a particular path. For example, C:\folder\log*.txt.

Excluding by file path

The **Folder Path** filter enables you to exclude files and folders in a specific path.

You can specify an absolute or relative path.

The following table provides examples for excluding files and folders using absolute and relative paths.

Table 12. Absolute and relative path examples

Type of path	Folder	File
Absolute	F:\folder1\folder2* In this example, the filter excludes all files and folders under F:\folder1\folder2.	F:\folder1\folder2\sample.txt In this example, the filter excludes the sample.txt file under F:\folder1\folder2.
Relative	*\folder1\folder2* In this example, the filter excludes all files and folders under any volume with the hierarchy folder1\folder2. D:*\folder1\folder2* In this example, the filter excludes all files and folders under any folder in D: with the hierarchy folder1\folder2.	*\folder1\folder2\sample.log In this example, the filter excludes all sample.log files under any volume with the hierarchy folder1\folder2. D:*\folder1\folder2\sample.log In this example, the filter excludes all sample.log files under any folder in D: with the hierarchy folder1\folder2.

Excluding multiple paths

Use commas to separate multiple paths and do not add a space between entries. For example, to exclude the folder C:\Program Files, the folder C:\Program Files (X86), and the folder C:\Perflogs, use the **Folder Path** filter C:\Program Files,C:\Program Files (X86),C:\Perflogs. Do not enclose each path in quotation marks.

Edit or delete an exclusion filter

Use the PowerProtect Data Manager to edit or delete an exclusion filter. You can change the filter name, description, logical operator, and filtering conditions.

Steps

1. Select **Protection > File Exclusion**.
The **File Exclusion** window appears with the added filters.
2. To edit a filter, complete the following tasks:
 - a. Select a filter, and click **Edit**.
The **Edit Filter** wizard appears.
 - b. Modify the desired fields, and then click **Next**.
The **Summary Confirmation** page appears.
 - c. Click **Finish** to save your changes.
3. To delete a filter, select the filter that you want to delete, and then click **Delete**.

Apply an exclusion filter to a protection policy

When adding or editing a protection policy, you can apply a predefined exclusion filter to the protection policy. The **File Exclusions** page of the **Add Policy** or **Edit Policy** wizard enables you to select an exclusion filter and apply it to a protection policy.

Prerequisites

An exclusion filter must exist.

About this task

To exclude assets within the protection policy from data protection operations, click **File Exclusions** when adding a policy.

To apply an exclusion filter to an existing protection policy, complete the following steps:


Steps

1. Select **Protection > Protection Policies**.
The **Protection Policy** window appears.
2. Select a protection policy from the list, and then click **Edit**.
The **Edit Policy** page appears.
3. Click **File Exclusions > Edit**.
The **File Exclusions** page appears.
4. Toggle the **Disabled** switch to enable exclusion.
5. Add a saved filter or build a new filter according to the steps provided in Add an exclusion filter.
6. Click **Next** twice, review the details on the **Summary** page, and click **Finish**.
PowerProtect Data Manager applies the exclusion filter to the protection policy.

Results

Once the backup job completes, you can view details about the number of files/folders that are excluded from the protection policy. To view the number of excluded files/folders:

1. In the PowerProtect Data Manager UI, select **Jobs > Protection Jobs** and click the **Job ID**. The **Protection Jobs > Job ID: <job ID>** page appears.
2. Click the **Details** icon. The **Step Log** and **Details** pane appear.
3. Under the **Details** pane, in the **Summary** section, you can find the number of files/folders that are excluded.

 **NOTE:** Only the debug logs contain file/folder names that are excluded from the backup.

Remove an exclusion filter from a protection policy

The **File Exclusions** page of the **Edit Policy** wizard enables you to remove an exclusion filter from a protection policy.

Steps

1. Select **Protection > Protection Policies**.
The **Protection Policy** window appears.
2. Select a protection policy from the list, and then click **Edit**.
The **Edit Policy** page appears.
3. Select **File Exclusions > Edit**.
4. Clear the check box next to the filter that you want to remove from the protection policy.
5. Click **Next**.
The **Summary** page appears.
6. Review the details, and click **Finish**.

Centralized restore of a file-system asset

When file systems are protected within a protection policy in PowerProtect Data Manager, you can recover the file system data using the centralized image-level or file-level restore functionality in the PowerProtect Data Manager UI.

Prerequisites for restore of file-system assets

Review the following requirements before performing centralized image-level or file-level restores of file system assets:

- Both the PowerProtect Data Manager server and client must be at a minimum version of PowerProtect Data Manager 19.3.
- Ensure that the File System agent is not installed and running on the target volume.
- Ensure that there is sufficient space on the target volume for the restore.
- Ensure that the target or destination volume is not read-only.
- Cross-platform support for centralized file-level restore is not supported. For example, you cannot recover a Windows backup on a Linux platform and the opposite way.
- Review the section Supported platform and OS versions for file system file-level restore.

Caution regarding image-level restore of a system volume to a system volume

Running an image-level restore from the backup of a system volume to a target volume that is the same or different system volume can cause the following problems:

- Files in use are not restored.
- The file system host machine may become unstable.

Therefore, in such a situation, it is recommended to perform a file-level restore for the required files and folders only.

Increasing the restore timeout

By default, the mount operation times out after 30 minutes and the backup copy is unmounted. When running file-level restores of large files, you can increase the restore timeout. Perform this task if file-level restores for large files timeout before completing.

1. Create a file with the name `browsersvc.cmd` in one of the following locations:
 - On Windows, `C:\Program Files\DPSAPPS\fsagent\settings`
 - On Linux, `/opt/dpsapps/fsagent/settings`
2. Add the following line to the file, and specify the same timeout value, in minutes, for both variables:

```
{"-idletimeout": "timeout", "-resexpiry": "timeout"}
```

For example, enter this line to set the restore timeout to 60 minutes:

```
{"-idletimeout": "60", "-resexpiry": "60"}
```

Centralized image-level restore of a file-system asset



A file-system host image-level restore enables you to recover data from backups of file systems performed in the PowerProtect Data Manager UI.

Prerequisites

- On Linux, before performing an image-level restore of the block-based backup copy, ensure that you are not logged in to the destination file system asset (volume) for other data protection operations. If you are logged in to the destination asset (volume) for any other data protection operation, the restore fails.

Steps

1. From the PowerProtect Data Manager UI, select **Restore > Assets**, and then select the **File System** tab.


The **Restore** window displays all of the file systems available for restore. Use   to switch between a list view of all file system assets and a Hierarchical view of assets within each File System host that has been discovered in PowerProtect Data Manager.

2. Select the check box next to the desired file system, and then click **View Copies**.

You can also use the **Search** field or the filter in the **Name** column to search on specific criteria to locate a specific file system.

The **Recovery > Assets** window provides a map view in the left pane and copy details in the right pane.

When you select a file system in the map view, the file system name appears in the right pane with the copy locations underneath. When you select a specific location in the left pane to view the copies, for example, on a DD system, the copies on that system appear in the right pane.

3. Click , and then select one of the available copies that appear in the table.
4. In the right pane, select the check box next to the file system backup you want to restore, and then click **Restore**. The **Restore** wizard appears.
5. On the **Select Target Location** page, choose from one of the following options, and then click **Next**.

- **Restore to original**—Restore the file system to the original location.
- **Restore to a new location on the original host**—Select the destination file system asset (volume) from the list of available assets on the host.
- **Restore to a new host**—Browse from the available hosts to locate and select a destination host and file system.

 **NOTE:** If the destination file system asset already contains some data, this data will be overwritten.

6. You can select **Troubleshooting mode** to enable debug logging, and then select the level of logging to use:
 - **Info**—Includes information such as status changes. This is the default log level for scheduled backups and restores.
 - **Debug**—Additional information that helps with problem diagnosis.
 - **Trace**—The most detailed amount of information for complex problem diagnosis.

The **Summary** page appears.

7. On the **Summary** page:
 - a. Review the information to ensure that the restore details are correct.
 - b. Click **Restore**.
8. Go to the **Jobs** window to monitor the restore.
A restore job appears with a progress bar and start time.

Centralized file-level restore of a file-system asset

A file-level restore enables the administrator to recover individual files from backups of file systems that were created in PowerProtect Data Manager.

Prerequisites

- Ensure the following for Linux file-system hosts:
 - You have enabled the SELinux `nfs_enabled` parameter by running one of the following relevant commands:
 - RHEL 8.x or CentOS 8.x: `setsebool -P nfs_enabled 1`
 - RHEL 7.x or CentOS 7.x: `setsebool -P nfs_enabled 1`
 - RHEL 6.x or CentOS 6.x: `setsebool -P allow_yphind 1`
 - You can also disable SELinux permanently:
 1. Open the `/etc/sysconfig/selinux` file in a text editor.
 2. Change the value of `SELINUX=enforcing` from `enforcing` to `disabled`.
 3. Restart the host machine.
 4. Verify the SELinux status by running the `getenforce` command.
 - You have installed the `iscsiadm` utility by installing one of the following relevant packages on the Linux client:
 - RHEL or CentOS: `iscsi-initiator-utils<version_number>.rpm`
 - SLES: `open-iscsi<version_number>.rpm`
 - On SLES, if you want to start the `iscsiadm` utility for the first time, restart the iSCSI services by running the following command: `service open-iscsi restart`
 - Review the section Supported platform and OS versions for file system file-level restore for supported platform and operating system versions. PowerProtect Data Manager supports file-level restore only if the backup or replica is on a DD system device.
- Ensure that you increase the restore request timeout from 10 to 120. This enables you to successfully complete the restoration process for a large number of files. Perform the following steps to increase the restore request timeout.
 1. Log in to the SSH of PowerProtect Data Manager as an admin.
 2. Run the following command to navigate to the `/usr/local/brs/lib/adm/config` directory.


```
cd /usr/local/brs/lib/adm/config
```
 3. Copy the `application.yml` file by running the following command.

```
cp application.yml application.yml.orig
```
 4. Open the `application.yml` file by running the following command.

```
vi application.yml
```
 5. In the `application.yml` file, change the `requestTimeout` from 10 to 120.
 6. Save the `application.yml` file and restart PowerProtect Data Manager.


Steps

1. From the PowerProtect Data Manager UI, go to **Recovery > Assets**, and then select the **File Systems** tab.

The **Restore** window displays the file systems that are available for restore. Use    to switch between a list view of all file system assets and a hierarchical view of assets within each File System host that has been discovered in PowerProtect Data Manager.


2. Select the check box next to the file system, and then click **View Copies**.

You can also use the filter in the **Name** column to search for the name of the specific file system or click the **File Search** button to search on specific criteria. See File system file-level restore from a search for more information.

 **NOTE:** The **File Search** is not supported for Linux files.

The **Restore > Assets** window provides a map view in the left pane and copy details in the right pane.

When a file system is selected in the map view, the file system name appears in the right pane with the copy locations underneath. When you select a specific location in the left pane to view the copies, for example, on a DD system, the copies on that system appear in the right pane.

3. If the backup is on a DD system, click , and then select from one of the available copies that appear in the table.
4. In the right pane, select the check box next to the file system backup you want to restore, and then click **File Level Restore**.
The **File level restore** wizard appears.
5. On the **Select target host and mount** page, choose from one of the following options, and then click **Mount**.
 - **Restore to Original.**
 - **Restore to Alternate.**

i **NOTE:** By default, centralized file-level restore operations restore the entire path of selected files and folders to the destination folder. For example, if the file `D:\Folder\file1.txt` is restored to `G:\CFLR`, it is restored as `G:\CFLR\D\Folder\file1.txt` instead of `G:\CFLR\file1.txt`. However, if you select the **Do not retain the upward folder hierarchy while performing recovery** option, only the selected file will be restored not the entire folder hierarchy. For example, if the file `D:\Folder\file1.txt` is restored to `G:\CFLR`, it is restored as `G:\CFLR\file1.txt`.
6. When the mount is complete, click **Next**.
The **Select folder and files to recover** page appears.
7. On the **Select folder and files to recover** page:
 - a. Expand individual folders to browse the original file system backup, and select the objects that you want to restore to the destination file system.
You can also use the filter in the **Name** column to search for the name of the specific object.
 - b. Click **Next**.
The **Select restore location** page appears.

i **NOTE:** When restoring a larger number files without increasing the request timeout, the **No files are available for restore** error appears on the **Select folder and files to recover** page. See the prerequisites of this topic for the procedure to increase the restore request timeout.
8. On the **Select restore location** page:
 - a. Select the destination drive. Alternatively, choose the **Overwrite files in restore location** option, in which case existing files on the destination drive will be overwritten.

i **NOTE:**
If you choose not to overwrite files and the file or folder has the same name as an existing file or folder, the selected file is renamed either before or after the file extension:

 - On Windows, the selected file or folder is renamed before the file extension. For example, `file1.txt` is renamed to `file1-SSID-timestamp.txt`.
 - On Linux, the selected file or folder is renamed after the file extension. For example, `file1.txt` is renamed to `file1.txt-SSID-timestamp`.
 - b. Browse the folder structure of the destination file system to select the folder where you want to restore the objects.
 - c. Click **Next**.
9. You can select **Troubleshooting mode** to enable debug logging, and then select the level of logging to use:
 - **Info**—Includes information such as status changes. This is the default log level for scheduled backups and restores.
 - **Debug**—Additional information that helps with problem diagnosis.
 - **Trace**—The most detailed amount of information for complex problem diagnosis.The **Summary** page appears.
10. On the **Summary** page:
 - a. Review the information to ensure that the restore details are correct.
 - b. Click **Finish**.
11. Go to the **Jobs** window to monitor the restore.
A restore job appears with a progress bar and start time.

File system file-level restore from a search

Within the **Restore** window of the PowerProtect Data Manager UI, click the **File Search** button. The **Search Criteria** pane displays to search files based on specific criteria. In the **Search Criteria** pane, enter or select any one or a combination of the following fields. The files that match the search criteria display in the **Results** pane.

Table 13. File search criteria

Search criteria	Description
File Name	Enter the complete or partial name of the file.
File Type	Enter the file type. For example, xls or doc .
Size	Enter the minimum or maximum size of the file and select a unit such as KB, MB, GB, or TB.
Folder Path	Enter the complete or partial folder path of the file.
Show Only Files	Move this slider to the right to filter only the files.
File System Name	Select or enter the file system name.
Host Address	Enter or select the name of the file system host.
Last Backup Only	Move this slider to the right to filter the files that are backed up recently.
Backup Date	Enter or select the specific date range to filter the files based on their backup date.
Date Modified	Enter or select the specific date range to filter the files based on their modified date.
Date Created	Enter or select the specific date range to filter the files based on their created date.

File Search enables you to restore files from protected file-system backup copies to:

- The original file-system host
- An alternate file-system host

NOTE:

- File-search functionality is supported for Windows NTFS file systems.
- File-level restore using **File Search** only supports restore of a single file or directory.
- For file-level restores, the files must be restored from a Windows backup to a Windows file system.

File-level restore to original file-system host using File Search

Use **File Search** in the PowerProtect Data Manager UI to restore files from one copy to the original file system host. Only the Administrator and the Restore Administrator roles can restore data.

Steps

1. From the PowerProtect Data Manager UI, select **Restore > Assets**, and then select the **File System** tab.
The **Restore** window displays all the file systems available for restore.
2. Click **File Search**, and then perform the following:
 - a. Select a file system from the **Name** list.
 - b. Use the **File Name** and **File Type** fields to search for specific files, or specify a file size or folder path to perform the search.
The files that match the search criteria display in the **Results** pane.
 - c. In the **Results** pane, select the file that you want to restore, and then click **Add**.
The **Results** pane is collapsed, and the **Selected Files** pane updates to display the current file selections.
 - d. When finished with your selections, click **Restore**.
The **File System Restore** wizard appears, displaying the **Location** page.
3. On the **Location** page:
 - a. Select **Restore to original host**.
 - b. Click **Next**.
The **Options** page appears.

- You can select **Troubleshooting mode** to enable debug logging, and then select the level of logging to use:
 - Info**—Includes information such as status changes. This is the default log level for scheduled backups and restores.
 - Debug**—Additional information that helps with problem diagnosis.
 - Trace**—The most detailed amount of information for complex problem diagnosis.

The **Summary** page appears.

- On the **Summary** page:
 - Review the information to ensure that the restore details are correct. You can click **Edit** next to certain rows to change the information.
 - Click **Restore** or **Finish**.
- Go to the **Jobs** window to monitor the restore.

A batch file-level restore job appears as a job group, with a progress bar and start time. A separate job entry is created for each copy that is being restored from.

File-level restore to alternate file-system host using File Search

Use **File Search** in the PowerProtect Data Manager UI to restore files from one copy to a new file system. Only the Administrator and the Restore Administrator roles can restore data.

Steps

- From the PowerProtect Data Manager UI, select **Restore > Assets**, and then select the **File System** tab.

The **Restore** window displays all the file systems available for restore.
 - Click **File Search**, and then perform the following:
 - Select a file system from the **Name** list.
 - Use the **File Name** and **File Type** fields to search for specific files, or specify a file size or folder path to perform the search.

The files that match the search criteria display in the **Results** pane.
 - In the **Results** pane, select the file that you want to restore, and then click **Add**.

The **Results** pane is collapsed, and the **Selected Files** pane updates to display the current file selections.
 - When finished with your selections, click **Restore**.

The **File System Restore** wizard appears, displaying the **Location** page.
 - On the **Location** page:
 - Select **Restore to alternate host**.


The table on the page updates to display the available destination file systems.
 - Select the file system.
 - Click **Next**.

The **Options** page appears.
 - You can select **Troubleshooting mode** to enable debug logging, and then select the level of logging to use:
 - Info**—Includes information such as status changes. This is the default log level for scheduled backups and restores.
 - Debug**—Additional information that helps with problem diagnosis.
 - Trace**—The most detailed amount of information for complex problem diagnosis.
- The **Summary** page appears.
- On the **Summary** page:
 - Review the information to ensure that the restore details are correct. You can click **Edit** next to certain rows to change the information.
 - Click **Restore** or **Finish**.
 - Go to the **Jobs** window to monitor the restore.

A batch file-level restore job appears as a job group, with a progress bar and start time. A separate job entry is created for each copy that is being restored from.

Supported platform and OS versions for file system file-level restore

File system file-level restore is only supported for the following platforms and operating systems.

-  **NOTE:** Platforms and operating systems are qualified for file-level restore support using the default file system for these platforms:

- RedHat Enterprise Linux
- SuSE Linux Enterprise Server
- CentOS
- Windows

Linux platforms require an ext3, ext4, XFS, or BTRFS file system type.

i **NOTE:** Refer to the E-Lab Navigator for the most up-to-date software compatibility information for PowerProtect Data Manager software and the File System agent.

Updating the File System agent hostname or IP address

You can change the hostname or IP address of the File System agent host.

To ensure assets continue to be protected, follow either Update the File System agent hostname or Update the File System agent IP address, or reregister the agent after an agent host or PowerProtect Data Manager IP address change.

Update the File System agent hostname

Perform the following steps to update the File System agent hostname.

Steps

1. Delete the asset source from PowerProtect Data Manager. This automatically deletes assets from the host that are added in the protection policies. See Delete an asset source for more information.
2. Change the hostname of the File System agent.
3. Delete the existing `agents.clb*` lockbox files in the `C:\Program Files\DPSAPPS\common\lockbox` directory on Windows or the `/opt/dpsapps/fsagent/lockbox` directory on Linux.

i **NOTE:** If the File System agent is installed to a nondefault path, delete all the files in the `lockbox` subdirectory of the installation directory.

4. Reregister the agent. To reregister the agent for Windows, run the installer with the **Change** option. To reregister the agent for Linux or AIX, run `register.sh --enable` from the `agentsvc` directory. This reregisters the deleted File System agent service. For more information, see Reregister the File System agent on Linux or AIX and Reregister the File System agent on Windows.
5. Once the client is registered with the new hostname, add the assets to the existing protection policy to perform backups. If there no policy exists, create a policy and add the assets to that.

Update the File System agent IP address, or reregister the agent after an agent host or PowerProtect Data Manager IP address change

Perform the following steps to update the File System agent IP address. These steps must also be followed when either the agent host IP address has already been changed or the PowerProtect Data Manager IP address has changed.

Steps

1. Unregister the agent. To unregister the agent for Windows, run `unregister.bat` from the `<Install_folder>\AgentService` folder. To unregister the agent on Linux or AIX, run `unregister.sh` from the `agentsvc` directory.

i **NOTE:** Do not follow this step if the agent host IP address has already been changed or if the PowerProtect Data Manager IP address has changed.

Change the IP address of the File System agent host.

3. Reregister the agent. To reregister the agent for Windows, run `register.bat` from the `<Install_folder>\AgentService` folder. To reregister the agent for Linux or AIX, run `register.sh` from the `agentsvc` directory.

Manage the PowerProtect agent service

The PowerProtect agent service provides important functionality for the application agent operations with the PowerProtect Data Manager.

Review the following topics to ensure that you enable and manage the PowerProtect agent service functionality as required for application agent operations.

About the PowerProtect agent service

The PowerProtect agent service is a REST API based service that is installed by the application agent on the application host. The agent service provides services and APIs for discovery, protection, restore, instant access, and other related operations. The PowerProtect Data Manager uses the agent service to provide integrated data protection for the application assets.

This section uses `<agent_service_installation_location>` to represent the PowerProtect agent service installation directory. By default, the agent service installation location is `C:\Program Files\DPSAPPS\AgentService` on Windows and `/opt/dpsapps/agentsvc` on Linux. All files that are referenced in this section are the relative paths to the agent service installation location.

The PowerProtect agent service performs the following operations:

- Addon detection—An addon integrates the application agent into the agent service. The agent service automatically detects the addons on the system for each application asset type and notifies the PowerProtect Data Manager. While multiple addons can operate with different asset types, only one agent service runs on the application host. Specific asset types can coexist on the same application host.
 - Discovery—The agent service discovers both stand-alone and clustered file system assets on PowerProtect Data Manager hosts, as well as their backup copies. After an initial discovery when the agent service discovers new assets and backup copies, the agent service notifies PowerProtect Data Manager.
 - Self-service configuration—The agent service can configure the application agent for self-service operations by using information that is provided by the PowerProtect Data Manager. When you add an asset to a protection policy for self-service or centralized protection, or modify the protection policy, including changing the DD Boost credentials, the PowerProtect Data Manager automatically pushes the protection configuration to the agents.
 - Centralized backups—The agent service performs the centralized backups as requested by the PowerProtect Data Manager.
 - Centralized restores—The agent service performs the centralized restores as requested by the PowerProtect Data Manager.
- NOTE:** In the current release, the centralized restores are only available for the File System agent, Microsoft SQL agent, and Storage Direct agent.
- Backup deletion and catalog cleanup—The PowerProtect Data Manager deletes the backup files directly from the protection storage when a backup expires or an explicit delete request is received. The PowerProtect Data Manager goes through the agent service to delete the catalog entries from the agent's local datastore.

- NOTE:** The manual deletion of backup copies is not recommended. PowerProtect Data Manager automatically deletes expired backup copies as needed.

The agent service is started during the agent installation by the installer. The agent service runs in the background as a service and you do not interact with it directly.

The `config.yml` file contains the configuration information for the agent service, including several parameter settings that you can change within the file. The `config.yml` file is located in the `<agent_service_installation_location>` directory.

If the `config.yml` file becomes corrupted, you can run the following commands to restore the file and continue the protection provided by the agent service:

- On Windows:

```
agentService.exe config=config.yml service=false restoreConfig=true
```

- On Linux and AIX:

```
agentService config=config.yml service=false restoreConfig=true
```

The agent service periodically starts subprocesses to perform the discovery jobs. You can see the type and frequency of these jobs in the `jobs:` section of the `config.yml` file. The job interval unit is minutes.

The agent service maintains a datastore in the `<agent_service_installation_location>/dbs/v1` directory, which contains information about the application system, assets, and backups discovered on the system. The size of the datastore files depends on the number of applications and copies on the host. The agent service periodically creates a backup of its datastore in the `<agent_service_installation_location>/dbs/v1/backups` directory, as used to recover the datastore if this datastore is lost.

NOTE: The size of each datastore backup is the same as the datastore itself. By default, a backup is created every hour. To save space on the file system, you can reduce this datastore backup frequency for large datastores. By default, the datastore backup is retained for one week. You can change the datastore backup frequency, retention period, and backup location in the `config.yml` file.

Changing the preferred host address

You can change the preferred host address that is used for communication between PowerProtect Data Manager and the application-agent host.

The preferred host address can be an IPv4 address, an IPv6 address, or a fully qualified domain name (FQDN).

NOTE: Using an FQDN is only available as of version 19.14 of the application agent.

It is recommended to use an FQDN as the preferred host address in DHCP environments and other environments where static IP addresses can change. When DHCP is used instead of an IP address and the IP address of the host changes, the following benefits occur:

- Communication between PowerProtect Data Manager and the host is uninterrupted.
- Backup and restore operations succeed.
- Reregistration of the application agent is not required.

Change the preferred host address

About this task

To change the preferred host address, perform the following steps:

Steps

1. From the left navigation pane, select **Infrastructure > Application Agents**.
2. Select the entry for the host.
3. Click **More Actions** and select **Set Preferred Address**.
4. From the **Preferred Address** drop-down, select the preferred address.

NOTE: DNS name resolution must be enabled to select a fully qualified domain name. If it is not enabled, enable it by selecting **Infrastructure > Application Agents** and clicking **Configure DNS Name Resolution**. After it is enabled, repeat these steps.

Start, stop, or obtain the status of the PowerProtect agent service

The PowerProtect agent service is started during the agent installation by the installer. If needed, you can use the appropriate procedure to start, stop, or obtain the status of the agent service.

On AIX or Linux, you can start, stop, or obtain the status of the agent service by running the `register.sh` script that is found in the `<agent_service_installation_location>` directory.

- To start the agent service:

```
# register.sh --start
Started agent service with PID - 1234
```

Alternatively on Linux, you can use the following command to start the agent service:

```
# service agentsvc start
```

- To stop the agent service:

```
# register.sh --stop
Successfully stopped agent-service.
```

Alternatively on Linux, you can use the following command to stop the agent service:

```
# service agentsvc stop
```

- To obtain the status when the agent service is running:

```
# register.sh --status
Agent-service is running with PID - 1234
```

- To obtain the status when the agent service is not running:

```
# register.sh --status
Agent-service is not running.
```

- Alternatively on Linux, you can use the following command to obtain the status of the agent service when it is running or not running:

```
# service agentsvc status
```

Recovering the PowerProtect agent service from a disaster

You can perform self-service restores of application assets by using a file system or application agent, regardless of the state of the agent service or PowerProtect Data Manager. The information in this section describes how to bring the agent service to an operational state to continue if a disaster occurs and the agent service datastore is lost.

The agent service periodically creates a backup of its datastore in the `<agent_service_installation_location>/dbs/v1/backups` repository. If all of these backups are lost, the agent service can still start. The agent service discovers all the application systems, assets, and backup copies on the system again, and notifies PowerProtect Data Manager. Depending on when the failure occurred, the agent service might not be able to find older backup copies for some asset types. As a result, the centralized deletion operations might fail when cleaning up the database vendor catalog or removing older backups that are taken before the asset is added to PowerProtect Data Manager.

By default, the agent service backs up consistent copies of its datastore files to the local disk every hour and keeps the copies for 7 days. Each time the agent service backs up the contents of the datastore, it creates a subdirectory under the `<agent_service_installation_location>/dbs/v1/backups` repository. The subdirectories are named after the time the operation occurred, in the format `YYYY-MM-DD_HH-MM-SS_epochTime`.

By default, the datastore repository is on the local disk. To ensure that the agent service datastore and its local backups are not lost, it is recommended that you back up the datastore through file system backups. You can also change the datastore backup location to a different location that is not local to the system. To change the datastore backup location, update the values in the `config.yml` file.

Restore the PowerProtect Data Manager agent service datastore

Prerequisites

NOTE: Ensure that the agent service is powered off. Do not start the agent service until disaster recovery is complete.

About this task

You can restore the datastore from the datastore backup repository. If the repository is no longer on the local disk, restore the datastore from file system backups first.

To restore the datastore from a backup in the datastore backup repository, complete the following steps:

Steps

1. Move the files in the `<agent_service_installation_location>/dbs/v1` directory to a location for safe keeping.

NOTE: Do not move or delete any `<agent_service_installation_location>/dbs/v1` subdirectories.

2. Select the most recent datastore backup.

The directories in the datastore backup repository are named after the time the backup was created.

3. Copy the contents of the datastore backup directory to the `<agent_service_installation_location>/dbs/v1` directory.

After the copy operation is complete, the `<agent_service_installation_location>/dbs/v1` directory should contain the following files:

- `copies.db`
- `objects.db`
- `resources.db`
- `sessions.db`

4. Start the agent service.

Performing Self-Service Backups and Restores with the File System Agent

Topics:

- Performing self-service backups of file systems
- Performing self-service restore of a file system host

Performing self-service backups of file systems

A host with the File System agent installed requires a PowerProtect Data Manager server to back up file systems.

To back up file systems manually and use PowerProtect Data Manager, register the host to PowerProtect Data Manager, create a self-service protection policy, and configure the retention policy.

NOTE: Select **Self-Service Protection** when you create the file systems protection policy in the PowerProtect Data Manager UI.

After a host is registered with PowerProtect Data Manager and assets are added to a self-service protection policy, use the `ddfsav` command to run self-service or manual backups on the host file system assets, as in the following example:

```
ddfsav -l FULL -a DFA_SI_DD_HOST=server_name -a DFA_SI_DD_USER=username -a
DFA_SI_DEVICE_PATH=storage_unit_and_path volume_names
```

where:

`-l {FULL | INCR}`

Specifies the type of the backup to perform such as full (FULL), or incremental (INCR). The default value is FULL.

`-a "DFA_SI_DD_HOST=server_name"`

Specifies the IPv4 address for the DD that contains the storage unit to back up the file system assets.

`-a "DFA_SI_DD_USER=username"`

Specifies the protection storage unit username. Example: **Policy-Protection**

`-a "DFA_SI_DEVICE_PATH=storage_unit_and_path"`

Specifies the name and the path of the storage unit where you want to direct the backup. Example: /**PolicyProtection/LVMs/2**

`volume_names`

Specifies one or more file system volumes to be backed up. Example: **F:\ E:\ G:**

For more information about how to use the `admin` utility to query the list of backups for an asset, see *Using the ddfsadmin utility for file systems*.

To perform a self-service backup, use the storage unit and username that was created on the DD system when the policy was created. PowerProtect Data Manager discovers these backups and enables centralized restore operations. You can also perform a manual restore operation.

Performing self-service restore of a file system host

When file systems are protected within a protection policy in PowerProtect Data Manager, you can recover the file system data by using the centralized PowerProtect Data Manager restore functionality, or directly by using the self-service restore feature. The following section describes the procedure for self-service restore of file systems.

Prerequisites for file system restores

Before performing centralized or self-service file system restores:

- Ensure that the target or destination volume is not a system volume.
- Ensure that the **File System agent** is not installed and running on the target volume.
- Ensure that there is sufficient space on the target volume for the restore.

Using the ddfsadmin utility for file systems

The ddfsadmin utility provides the following command line options for file system recovery.

ddfsadmin backup query

Before running the `ddfsrhc` command to perform a self-service image-level restore of file systems, you can use the `ddfsadmin backup query` command to query a list of all the local and remote backups taken for a particular host, as shown in the following:

`ddfsadmin backup query -local -v=volume name -t=time value [h = hour,d = days,w = weeks,m = months]` queries the local record file for listing backups.

`ddfsadmin backup query -remote -d=Protection storage system -s=storage unit -u=username -p=DD password -c=hostname -v=volume name -t=time value [h = hour,d = days,w = weeks,m = months]` queries the record file on the protection storage system for listing backups.

Example usage

`ddfsadmin backup query -local -v="C:\\" -t=5` displays a list of local backups in C:\ taken within the last five days.

ddfsadmin sync

This command ensures that the catalogs that are on the local machine and in the DD system are synchronized. The following is the usage for the `ddfsadmin sync` command:

```
sync -local options: Sync local record file with record file on DD
sync -remote options: Sync remote record file with file in the local
options:
  -d=<DD host>: Protection storage system host IP
  -u=<DD username>: Protection storage system username
  -s=<DD device path>: Protection storage system device path
  -p=<DD password>: Protection storage system password.[Optional]
```

Example usage

`ddfsadmin sync -local -d x.x.x.x -u username -s /dev_path`

Self-service image-level restore of file systems

You can perform self-service image-level restores of file systems to the original location by using the `ddfsrhc` command. This restore is not supported in the following scenarios:

- When the restore destination is the C:\ volume, which can result in the operating system becoming unavailable.
- When the restore destination is a volume with the File System agent installed.

NOTE: To perform file system restore to an alternate location, use the centralized restore method in the PowerProtect Data Manager UI, as described in the section Centralized image-level restore of a file-system asset.

Before running `ddfsrc`, use the `ddfadmin backup` command to list the local backups for a particular host and obtain the ID of the save set you want to restore. Using the `ddfadmin` utility for file systems provides more information about the `ddfadmin backup` command.

To restore from a particular backup, specify the ID of the save set as an input to the `ddfsrc` command, as in this example:

```
ddfsrc -h DFA_SI_DEVICE_PATH=device path (for example, /fsa2) -h DFA_SI_DD_HOST=Protection storage system IPv4 address -h DFA_SI_DD_USER=Protection storage system username (for example, sysadmin) -S 1551407738 -r file path (for example, /volume1_ext3) -i y.
```

where:

`-h "DFA_SI_DEVICE_PATH=<storage_unit_and_path>"`

Specifies the name and the path of the storage unit that contains the backup.

`-h "DFA_SI_DD_HOST=<server_name>"`

Specifies the name of the protection storage system server that contains the backup.

When you have a remote (secondary) protection storage system server that has replicated databases to restore, type the name of the secondary server. A user on the secondary protection storage system server must be in the same group as the primary protection storage system server.

`-h "DFA_SI_DD_USER=<Protection storage system user>"`

Specifies the protection storage system username.

You must register the hostname and the DD Boost username in the lockbox to enable Microsoft application agent to retrieve the password for the registered user.

Self-service file-level restore of file systems

You can perform self-service file-level restores of file systems using the `ddfsrc` command with the `-I` option.

Before starting the command, create a file that contains the list of files to be restored. Provide the location of this file as an input to the `-I` option, as shown in the following example.

ddfsrc command with input file specified

```
ddfsrc -h DFA_SI_DEVICE_PATH=Protection storage unit -h DFA_SI_DD_HOST=Protection storage system IP address -h DFA_SI_DD_USER=Protection storage system username -S savetime-value -I path-of-file-containing-list-of-files-for-restore -i R -d destination-path-for-restoring-files
```

NOTE: In this command,

- The parameter `R` renames the file.
- If you replace the parameter `R` with `Y`, the file is overwritten.
- If you use `-g` at the end of the command, only the specified file is restored. Else, the entire directory structure is restored to the destination folder. For example, if the file `D:\Folder\file1.txt` is restored to `G:\CFLR`,
 - Without `-g`, the file is restored to `G:\CFLR\D\Folder\file1.txt`.
 - With `-g`, the file is restored to `G:\CFLR\file1.txt`.

The following steps provide more detail:

1. Use the `ddfadmin` command to list all the available backups. If you know the save set ID of the backup from which you want to restore, skip this step.

For example, the following command lists all backups taken in the last 55 days.

```
[root@XXXX ~]# ddfadmin backup query -local -t=55d
```

2. Create an input file that contains the list of files to restore. For example:

```
[root@XXXX ~]# cat flr.txt  
/new_ext3/file.txt
```

The `flr.txt` file specifies a single file to restore (`file.txt`).

3. Run the `ddfsrc` command. Ensure that you provide the complete path to the input file that you created.

① **NOTE:** Do not provide a relative path. If you provide a relative path, the command fails.

For example:

```
ddfsrc -h DFA_SI_DEVICE_PATH=Protection storage unit -h DFA_SI_DD_HOST=Protection storage system IP address -h DFA_SI_DD_USER=Protection storage system username -S savetime-value -I /root/flr.txt -d destination-path-for-restoring-files
```

where savetime-value is the save set ID identified in step 1.

Performing Self-Service Backups and Restores for Disaster Recovery

Topics:

- Performing self-service backups for disaster recovery
- Using the `ddfssv` utility for disaster recovery
- Self-service system state restore for disaster recovery
- Bare-metal recovery restore for disaster recovery

Performing self-service backups for disaster recovery

A host with the File System agent installed requires a PowerProtect Data Manager server for disaster recovery asset.

To back up file systems manually and use PowerProtect Data Manager, register the host to PowerProtect Data Manager, create a self-service protection policy, and configure the retention policy.

NOTE: Select **Self-Service Protection** when you create the file systems protection policy in the PowerProtect Data Manager UI.

After a host is registered with PowerProtect Data Manager and assets are added to a self-service protection policy, use the `ddfssv` command to run self-service manual backups on the host disaster recovery asset, as in the following example:

```
ddfssv -l FULL -a DFA_SI_DD_HOST=server_name -a DFA_SI_DD_USER=username -a
DFA_SI_DEVICE_PATH=storage_unit_and_path volume_names
```

where:

`-l {FULL | INCR}`

Specifies the type of the backup to perform such as full (FULL), or incremental (INCR). The default value is FULL.

`-a "DFA_SI_DD_HOST=server_name"`

Specifies the IPv4 address for the DD that contains the storage unit to back up the file system assets.

`-a "DFA_SI_DD_USER=username"`

Specifies the protection storage unit username. Example: **Policy-Protection**

`-a "DFA_SI_DEVICE_PATH=storage_unit_and_path"`

Specifies the name and the path of the storage unit where you want to direct the backup.

Example: `/plc_self_service_protectionpolicy/PLCTLP-ab31adac-1a4f-4d26-9d00-a3148d63805a`

`-a "DFA_SI_DR_SSRONLY=TRUE"`

Specifies whether the SSR needs to be backed up or not. Set the flag value to **TRUE** if you want to perform SSR backup. The default value of this flag is **FALSE**. This is optional for BMR backups.

`-a "DFA_SI_IGNORE_MISSING_VSS_FILES= TRUE"`

Specifies whether the missing windows system state files must be ignored during the self-service backup. The default value of this flag is **TRUE**. This flag states that the backup is completed successfully instead of stating the backup is completed with exceptions. However, the corresponding logs list the missing files as warnings.

NOTE: If the flag value of this command is set to **FALSE**, the missing files will be listed as errors in logs and the backup is completed with exception.

`volume_names`

Specifies Disaster Recovery Asset to be backed up. Example: `DISASTER_RECOVERY ; \`

For more information about how to use the `admin` utility to query the list of backups for an asset, see [Using the `ddfsadmin` utility for disaster recovery](#).

To perform a self-service backup, use the storage unit and username that was created on the DD system when the policy was created. PowerProtect Data Manager discovers these backups and enables centralized restore operations. You can also perform a manual restore operation.

Using the `ddfsadmin` utility for disaster recovery

The `ddfsadmin` utility provides the following command line options for disaster recovery asset.

`ddfsadmin backup query`

Before running the `ddfsrc` command to perform a self-service system state restore for disaster recovery asset, you can use the `ddfsadmin backup query` command to query a list of all the local and remote backups taken for a particular host, as shown in the following:

`ddfsadmin backup query -local -t=time value [h = hour, d = days, w = weeks, m = months]` queries the local record file for listing backups.

Example usage

`ddfsadmin backup query -local -t=5d` displays a list of local backups in `DISASTER_RECOVERY:\` taken within the last five days.

`ddfsadmin sync`

This command ensures that the catalogs that are on the local machine and in the DD system are synchronized. The following is the usage for the `ddfsadmin sync` command:

```
ddfsadmin sync -local -d x.x.x.x -u username -s /dev_path
```

options:

```
-d=<DD host>: Protection storage system host IP  
-u=<DD username>: Protection storage system username  
-s=<DD device path>: Protection storage system device path  
-p=<DD password>: Protection storage system password. {Optional}
```

Self-service system state restore for disaster recovery

You can perform self-service system state restores for disaster recovery asset using the `ddfsrc` command with the `-I` option.

Before starting the command, create a file that contains the list of writers to be restored. Provide the location of this file as an input to the `-I` option, as shown in the following example.

`ddfsrc` command with input file specified

```
ddfsrc -h DFA_SI_DD_HOST=Protection storage system IP address -h DFA_SI_DD_USER=Protection storage system username -h DFA_SI_DEVICE_PATH=Protection storage unit -h DFA_SI_DR_SSRONLY=TRUE -I path-of-file-containing-list-of-writers-to-restore -S savetime-value
```

where:

```
-a "DFA_SI_DD_HOST=server_name"
```

Specifies the name of the protection storage system server that contains the backup. When you have a remote (secondary) protection storage system server that has replicated databases to restore, type the name of the secondary server. A user on the secondary protection storage system server must be in the same group as the primary protection storage system server.

-a "DFA_SI_DD_USER=protection_storage_system_user"

Specifies the protection storage system username. You must register the hostname and the DD Boost username in the lockbox to enable Microsoft application agent to retrieve the password for the registered user.

-a "DFA_SI_DEVICE_PATH= storage_unit_and_path"

Specifies the name and the path of the storage unit that contains the backup.

-a " DFA_SI_DR_SSRONLY=TRUE"

Specifies the flag status to TRUE if you want to perform SSR restore. The default flag value is FALSE.

-I "path-of-file-containing-list-of-writers-to-restore"

Specifies the path of the file containing the list of writers to be restored.

-S "savetime-value"

Specifies the save set ID of a backup copy which needs to be restored.

The following steps provide more detail:

1. Use the `ddfsadmin` command to list all the available backups. If you know the save set ID of the backup from which you want to restore, skip this step.

For example, the following command lists all backups that are taken in the last 55 days.

```
ddfsadmin backup query -local -t=55d
```

2. Create an input file that contains the list of writers to restore. For example:

```
Notepad.exe ssr.txt
```

- The `ssr.txt` file specifies a single writer to restore. For example, **System Writer**
or

- The `ssr.txt` file specifies multiple writers that must be restored, where each writer name should be valid and specified in a new line. For example,

```
System Writer
```

```
Registry Writer
```

```
Task Scheduler Writer
```

```
WMI Writer
```

NOTE: If the writer name is invalid, the restore operation will be skipped for that writer.

3. Run the `ddfsarc` command. Ensure that you provide the path to `ssr.txt` file that you created.

For example:

```
ddfsarc -h DFA_SI_DD_HOST=Protection storage system IP address -h  
DFA_SI_DD_USER=Protection storage system username -h DFA_SI_DEVICE_PATH=Protection  
storage unit -h DFA_SI_DR_SSRONLY=TRUE -I C:\\ssr.txt -S savetime-value
```

where `savetime-value` is the save set ID identified in step 1.

Bare-metal recovery restore for disaster recovery

You can perform a bare-metal recovery restore for disaster recovery of assets.

The E-Lab Navigator provides the most up-to-date information about the operating systems that support bare-metal recovery restores.

To perform a bare-metal recovery restore of Windows assets using WinPE, see [Perform a bare-metal recovery](#).

Performing Disaster Recovery with the File System Agent in Windows

Topics:

- Disaster recovery limitations
- Preparing for disaster recovery
- Performing system-state recovery
- Recovering the Active Directory
- Performing bare-metal recovery
- Perform bare-metal recovery of Windows clusters
- Disaster recovery solution for Microsoft SQL application data
- Protecting Microsoft Distributed File System using BMR and SSR
- Performing application restores after bare-metal recovery

Disaster recovery limitations

The following limitations apply to performing disaster recovery in the File System agent.

- Disaster-recovery data of clusters does not include critical disks in shared storage. Critical disks must be backed up separately.
- Physical host to virtual host (P2V) BMR of a Hyper-V server to a VMware virtual machine is not supported.
- By default, disaster recovery saveset backup is performed as file-based backups. Block-based backup is not supported for disaster recovery.

Preparing for disaster recovery

Gathering key information

Before starting a disaster recovery operation, you must gather the following information about the relevant systems:

- File system configuration
- Hard drive configuration
- Device driver information for bare-metal recoveries

Critical volumes in disaster recovery

Critical volumes included in disaster-recovery data are shown when selecting assets to be backed up.

The following volumes are included in disaster-recovery data:

- Any volume that contains operating-system files.
- Any volume that contains a third-party service.
- Any noncritical volume that has a critical volume mounted on it, or any noncritical volume that serves as a parent to a critical volume. In either case, both the parent volume and mounted volume are treated as critical.
- If any of the volumes on a dynamic disk is critical, all volumes on the dynamic disk are considered critical. This is a Microsoft requirement.

Discover the assets to back up

If some application assets are not discovered, you can perform an immediate full discovery of application asset sources by using the on-demand discovery feature in the PowerProtect Data Manager UI.

About this task

To initiate a full discovery of application asset sources, complete the following steps:

Steps

1. From the left pane, select **Infrastructure > Asset Sources**.
2. On the **File System** tab, select the agent on which the assets are to be discovered.
3. Click **Discover**.
4. Click **Yes** to continue.

Results


You can view the progress of the discovery from the **Jobs > System Jobs** page. When the job completes and the asset is discovered, the **Status** is **Available**.

Create a disaster recovery protection policy

About this task

A disaster recovery protection policy should contain objects to be backed up, which include critical volumes and system-state recovery files.

Steps

1. From the left pane, select **Protection > Protection Policies**.
2. Click **Add**.
3. In the **Name** field, type a name for the policy.
4. Ensure that the **File System** option is selected, and then click **Next**.
5. Click **Next**.
6. In the **Assets** pane, select the assets that the policy covers, and click **Next**.
7. If a disaster recovery object was selected in the previous step, leave the **File Exclusions** feature *Disabled*, and then click **Next**.
8. In the **Objectives** pane, click **Next**.
9. In the **Disaster recovery options** pane, select the options that you want applied to the policy, and then click **Next**.
 - **Back up system state files only** - Performs a backup of system state files only. By default, this check box is not selected and bare-metal recovery (BMR) data is backed up.
 **NOTE:** If the policy is configured with this option, you can only perform a system-state recovery (SSR), and the backed up data will only contain SSR information. BMR with WinPE is not possible.
 - **Ignore missing system state files** - Missing Windows system state files are reported as errors, and the backup fails, reporting the files as missing. This option is selected by default.
 - **Exclude non-critical dynamic disks** - If any volume of a dynamic disk pack is critical, all volumes in the dynamic disk pack are considered critical. By default, this option is not selected and noncritical dynamic disks are included in the backup data. To avoid the creation of large system state files, select this option to exclude noncritical dynamic disks from the backup data.
 - **Ignore third-party services when identifying critical volumes** - When a Windows service or application is installed on an otherwise noncritical disk, that disk is considered critical. By default, this option is not selected and the backup includes the disks on which a Windows service or application is installed. To avoid the creation of large system state files, select this option.
10. Click **Next**.
11. Click **Finish**.

Results

You can view the progress of the policy creation from the **Jobs > System Jobs** window.

If you use the **Edit Policy** wizard to add a disaster recovery asset to an existing protection policy, the **Disaster recovery** pane is shown, with options that are the same as the options described in step 9.

Synchronize all clocks

To ensure discovery of all disaster recovery assets, ensure that the clocks on both the host and PowerProtect Data Manager are synchronized to the local Network Time Protocol (NTP) server.

Manually run a disaster recovery policy

Steps

1. From the left pane, select **Protection > Protection Policies**.
2. Select the disaster recovery policy that you want to run.
3. Click **Protect Now**.
4. Click **Next**.
5. Select whether you want to back up all assets or a customized set.
6. Click **Next**.
7. In the **Select backup type** drop-down list, select the type of backup that you want.
8. In the **Retain for** fields, specify how long you want the backup kept.
9. Click **Next**.
10. (Optional) To change the assets selected to be backed up or the configuration of the backup data, click **Edit** and make the necessary changes.
11. Click **Protect Now**.

Results

You can view the progress of the backup on the **Jobs > Protection Jobs** page.

Performing system-state recovery

System-state recovery (SSR) is an online recovery that enables you to recover an online or powered-on machine that has lost its system files and registry. Perform an SSR when you want to restore certain selected operating system files from a known good backup to replace the corrupted or missing files. You can perform granular recoveries of selected writers from backed up bare-metal recovery (BMR) or SSR data. All the system files in the backup can be recovered only to their original location. Recovering the data to alternate hosts or locations is not supported.

Following are a few examples of files and components that are included in an SSR:

- Boot files
- COM+ class registration database
- Registry and IIS metadata
- Active Directory (NTDS)
- System volume (SYSVOL)

Perform a system-state recovery



If system files or registry entries are lost, you can recover the relevant writers and perform a system-state recovery (SSR).

Prerequisites

- Ensure that the host for which the SSR is to be performed is powered on.
- Ensure that the writers on the host are available and in a stable state.

- Ensure that all the services on the client are running exactly as they were running during the SSR/BMR backup. All the services must be running for an SSR restore to be successful. If services are corrupted, stopped, or not available on the client, perform the BMR restore.
- When performing an SSR, ensure that there is at least 50% of free space on the system disk.

Steps

1. In PowerProtect Data Manager, from the left pane select **Restore > Assets**.
2. Select the check box for the relevant client.
3. Click **View Copies** to view the copies that are backed up.
The copy map consists of the root node and its child nodes. The root node in the left pane represents an asset, and information about copy locations appears in the right pane. The child nodes represent storage systems.
4. Select  to display the copies on that storage system.
5. Select the desired backup copy, and then click **System State Restore**. After the backup copy is mounted successfully, a list of backed up writers is displayed.
6. The **Disaster Recovery** asset page selects all writers by default for an SSR. If you deselect individual writers, the following message appears:
 **NOTE:** You must select the entire system state to restore. Partial selection of system state restore is not recommended unless it is for Active Directory restore.
7. If you deselected one or more writers and the warning appeared, click **OK**.
8. To start the SSR, click **Finish**.
9. To see the status of the SSR, from the left pane select **Jobs > Protection Jobs**.
10. Wait for the SSR to complete.
11. Restart the host for which SSR was performed.

 **CAUTION:** Failing to restart the host can result in system instability.

Recovering the Active Directory

When you want to recover the Active Directory specifically, you must choose only NTDS writer and perform an Active Directory (AD) restore.

To recover the Active Directory, perform the following steps:

1. Configure the client. Configure the client to boot into Directory Services Restore Mode provides more information.
2. Recover the Active Directory. Recover the Active Directory from disaster-recovery data provides more information.
3. Perform an authoritative or nonauthoritative Microsoft restore based on user configuration. Authoritative and nonauthoritative Microsoft restores provides more information.

Configure the client to boot into Directory Services Restore Mode

Before you recover the Active Directory from disaster-recovery data, configure the client to boot into Directory Services Restore Mode (DSRM).

Steps

1. Run the `msoconfig` command. The System Configuration window appears.
2. On the **Boot** tab, select **Safe boot**, and then select **Active Directory repair**.
3. Click **OK**.
4. Restart the computer into Directory Services Restore Mode (DSRM).


Recover the Active Directory from disaster-recovery data

About this task

After you have configured the client to boot into Directory Services Restore Mode (DSRM), recover the Active Directory from disaster-recovery data.

Steps

1. Open **PowerProtect Data Manager**.
2. On the **Restore** tab, in the list of clients, select the client that you want to recover.
3. In the drop-down list for the host, select **Disaster Recovery**.
4. Click **View Copies** to view the backed-up copies.
5. Click  to display the copies.
6. Select the desired backup copy and click **System State Restore**.
7. After the copies are mounted, select the Windows NT Directory Services (NTDS) writer in the Disaster Recovery folder, and then click **Next**.
8. Click **Finish** to start the system-state recovery (SSR).
9. Wait for the Active Directory recovery to complete.
10. Restart the client after the recovery completes.

 **NOTE:** While performing SSR restore of all writers, if an NTDS writer is a part of the backup, restore the NTDS writer alone in the Directory Services Restore Mode (DSRM), restart the system in normal mode and then restore other remaining writers excluding NTDS writer.


Authoritative and nonauthoritative Microsoft restores

You can perform either a nonauthoritative or an authoritative Microsoft restore of Active Directory.

- Use a nonauthoritative restore when Active Directory replication partners can return a domain controller to a known state. You restore the domain controller from a backup. When you restart the domain controller after the restore, other domain controllers replicate changes made after the backup.
- Use an authoritative restore to return a domain controller to a known state as the master copy. The data from the restored domain controller replicates to other domain controllers. An authoritative restore also enables you to mark specific organizational units (OUs) so that Active Directory objects replicate to other domain controllers. Replication partners do not overwrite the replicated objects.

The following Microsoft TechNet articles provide details on an authoritative restore:

- "Performing Authoritative Restore of Active Directory Objects" provides general details on an authoritative restore.
- "Mark an Object or Objects as Authoritative" provides details on the command syntax for marking items for an authoritative restore.

 **NOTE:** Microsoft recommends using a nonauthoritative restore or reinstallation to restore a domain controller. The Microsoft TechNet article "Performing Nonauthoritative Restore of Active Directory Domain Services" provides information about reinstating a domain controller with a nonauthoritative restore.

You can choose whether to perform a nonauthoritative restore or an authoritative restore based on user configuration.

Perform a nonauthoritative Microsoft restore

After the Active Directory recovery completes, restart the client normally. Other domain controllers replicate changes to the client after the restart.

Perform an authoritative Microsoft restore

In an authoritative Microsoft restore, the data from the recovered domain controller replicates to other domain controllers.

Steps

1. Open a command-prompt window and run `ntdsutil` to mark objects for the authoritative restore.

The objects replicate to other domain controllers during the authoritative restore. In addition, replication partners do not overwrite the replicated objects.

You can mark a single user object, an entire user subtree, containers, or the entire database. You can use Microsoft **ADSIEdit** to display Distinguished Names for AD objects.

For example, the following series of commands marks a user with an OU of `CN=Test User,CN=Users,DC=svr1,DC=mydomain,DC=com` for an authoritative restore:

```
ntdsutil
  activate instance NTDS
    authoritative restore
    restore object
    "CN=Test User,CN=Users,DC=svr1,DC=mydomain,DC=com"
    quit
quit
```

The Microsoft documentation provides details on using the `ntdsutil` utility for an authoritative restore.

2. If you used Windows System Configuration to configure the system to boot into DSRM, use Windows System Configuration again and clear **Safe boot** to enable the system to boot normally.
3. Restart the client.

Performing bare-metal recovery

Bare-metal recovery (BMR) is used as part of a disaster recovery plan that provides protection when a machine cannot start and you must recover everything. Disaster situations include hardware failure and cyberattacks.

You can use BMR when your host is not available due to a hardware failure or it cannot start. Use BMR for either of the following reasons:

- You want to recover a computer in its entirety after a hardware failure that has been repaired.
- You want to recover data to a new computer after a hardware failure that cannot be repaired. The new computer does not have an operating system, and the OS files must also be recovered from the old computer.

By default, BMR data is System State enabled.

BMR data consists of the following:

- The operating system files and all data except user data on critical volumes
 - ① **NOTE:** Critical volumes include the boot volume, the system volume, and the volume that hosts system state data, such as Active Directory and application services.
- All system state information

BMR can be used for any of the following operations:

- Physical machine to physical machine (P2P)
- Physical machine to virtual machine (P2V)
- Virtual machine to virtual machine (V2V)

① **NOTE:** P2V BMR of a Hyper-V server to a VMware virtual machine is not supported.

To protect a Windows host entirely, it is recommended that you back up BMR data for critical volumes and separately back up regular assets that contain user data.

Bare-metal recovery requirements

Before you perform bare-metal recovery (BMR), verify that the environment meets the following requirements and that you have the necessary information:

- The hardware on the target host is operational.
- The hardware configuration on the target host is similar to the hardware configuration on the source host from which the BMR data was obtained. Any hardware, driver, or firmware differences between the target and source hosts can cause BMR to fail.
- The size of the disks on the target host is equal to or greater than the size of the disks on the source host. BMR fails to initialize and format a disk when the disk size on the target host is less than the disk size on the source host, even if the target system disk size is sufficient for the BMR data. After the BMR operation, some unallocated space might remain. You can extend the partition size after the BMR operation to use this extra space.
- There are at least as many disks on the target host as there were on the source host. The disk LUN numbering on the target host must match the disk LUN numbering on the source host.
- Both the source and target hosts use 64-bit Windows.
- Both the source and target hosts boot using BIOS or both boot using UEFI.
- For the BMR of a UEFI system, a drive letter is available.
- The source host to be recovered is turned off before the BMR is started.
- A custom WinPE image is available.
- You have the following information available:
 - The IP address and network name of the target host.
 - The network name or IP address of the PowerProtect Data Manager server to use for the BMR operation.
 - Account credentials for the Admin account on the PowerProtect Data Manager server.
 - The source hostname. To obtain the source hostname from the PowerProtect Data Manager UI, select **Infrastructure > Asset Sources**. Keep a note of the displayed source hostname. If an incorrect source hostname is provided in the BMR wizard during the BMR operation, the backup copies are not displayed.

About the WinPE image

WinPE enables you to boot with a minimal subset of Windows features, but still access network resources, disks, and other resources from a command prompt. The custom PowerProtect Data Manager WinPE image contains the NIC and disk drivers for the Windows versions that the WinPE image supports.

You can burn the WinPE image to a CD, DVD, or USB flash drive, and then boot the target host from that media.

When you boot with a customized WinPE image, the boot process automatically starts the **PowerProtect Data Manager Bare Metal Recovery Wizard**.

Using a custom WinPE image

PowerProtect Data Manager provides a custom WinPE image that enables you to recover a source host to a target host without installing an operating system. Because local disks are not in use by the operating system, the recovery process can replace files without conflict.

The custom PowerProtect Data Manager WinPE image is based on Windows PE 10.0, and contains the NIC and disk drivers for the Windows versions that the WinPE image supports.

If the custom WinPE image does not contain the drivers for the NIC or disk devices on the source host that you are recovering, you can perform one of the following tasks:

- Copy the drivers to a USB flash drive, and then connect the drive after booting with the custom PowerProtect Data Manager WinPE image.
- Create a WinPE image that includes the drivers, and boot from that image. For more information, see (Optional) Adding NIC or disk drivers to the WinPE ISO file.

The drivers must meet the following requirements:

- 64-bit.
- Do not require a restart during installation.

NOTE: The WinPE environment loads only in memory, and changes are not persistent across a restart. If a restart prompt appears, you might be able to ignore the prompt. Most NIC drivers are plug-and-play.

Use the custom WinPE image with BMR

About this task

Download, modify, and deploy the custom WinPE image for the BMR of a Windows target computer by completing the following procedure.

Steps

1. Download the custom WinPE image from the PowerProtect Data Manager server. For more information, see [Download the custom WinPE image from the PowerProtect Data Manager server](#).
2. If the WinPE ISO image does not contain the drivers for the NIC or disk devices on the source host that you are recovering, and you do not want to load the drivers from a separate disk during the BMR operation, and then add the drivers to the WinPE image. For more information, see [Add NIC or disk drivers to the custom WinPE image](#).
3. Use one of the following methods to deploy the WinPE image:
 - To boot the target host locally, burn the WinPE image to a CD, DVD, or USB flash drive.
 - To boot the target host over the network, copy the WinPE image to a Windows Deployment Services (WDS) server. For more information, see [Add the custom WinPE image to a Windows Deployment Services server](#).

Download the custom WinPE image from the PowerProtect Data Manager server

About this task

Complete the following procedure to download the custom WinPE image.

Steps

1. Open a web browser and type the following URL:
`https://<server>`
where <server> is the DNS name or IP address of the PowerProtect Data Manager server.
2. Select **Settings > Downloads**.
3. Select **WinPE**.
4. Click **Download** to download the .iso file of the custom WinPE image.
5. Download the file to a temporary folder.

Add NIC or disk drivers to the custom WinPE image

You can modify the custom WinPE image to add NIC or disk device drivers so you don't have to use a separate disk during the BMR operation.

About this task

If the custom WinPE image does not provide NIC or disk device drivers for the source host, you add them to the image.

NOTE: Modifying the image in any way other than adding NIC or disk device drivers is unsupported.

Steps

1. Open the .iso file of the WinPE image with a utility like UltraISO or MagicISO.
2. Create a folder for the drivers at the top level of the folder structure. The following example creates a `Drivers` folder.



Figure 2. WinPE folders

- Copy the NIC or disk device drivers to the folder.
If you have different source hosts that require different NIC or disk device drivers, you can create a subfolder for each device driver.
- Save the WinPE image with a different name.

Add the custom WinPE image to a Windows Deployment Services server

You can choose to add the custom WinPE image to a Windows Deployment Services (WDS) server to enable the target host to boot over the network. The Microsoft TechNet website provides detailed steps to configure and use WDS.

About this task

- NOTE:** Booting through WDS method from a WinPE image over the network and booting from an ISO image that is connected to the host are supported. The other network boot methods are unsupported.

Steps

- Configure the WDS server.
- Add the WinPE image to the boot menu.
- Ensure that PXE booting is enabled on the WDS target host.
- Boot the target host from the WinPE image over the network.

Perform a bare-metal recovery

About this task

Ensure that the hardware on the target host is operational and that the target host is similar in make, model, and hardware configuration to the source host to be recovered. Also, review the additional requirements in Bare-metal requirements.

- CAUTION:** If the source host to be recovered is powered on, power it down before starting the bare-metal recovery.

Steps

- Boot the target host with the custom WinPE image, either locally or over the network.
The **PowerProtect Data Manager Bare Metal Recovery Wizard** Welcome page is shown.
 - Specify the date, time, and time zone for the host, and then click **Next**.
If you are restoring to a host in a different time zone or if the system date and time are incorrect, you must change the default date and time.
- NOTE:** If you specify an invalid date or time, the wizard attempts to correct it. Verify that the corrected date and time are accurate.

3. Select the network interface for communication with PowerProtect Data Manager during the BMR operation. If the required NIC driver is not in the list, click **Load Driver** to browse to it.

NOTE: The driver must not require a restart. The WinPE environment loads only in memory, and changes are not persistent across a restart. If a restart prompt appears, you might be able to ignore the prompt. Most NIC drivers are plug-and-play.

4. Click **Next**.

The **Hostname and Network** tab opens.

5. In the **Host name** field, type the hostname of the target host.

6. In the **DNS domain** field, type the domain name of the target host.

If the host resides in a workgroup instead of a domain, leave the field blank.

7. Select the **IPv4** tab to configure the network to communicate with PowerProtect Data Manager during the BMR operation.

8. In the **TCP/IP Address** section, select the IP address to use:

- If host IP addresses are assigned automatically, then select **Obtain an IP address automatically (DHCP)**. The network must be configured to support DHCP.
- If host IP addresses are static, select **Use the following IP** address; and then enter the IP address and the IPv4 subnet mask.
- If PowerProtect Data Manager is on a different subnet, and then type the default gateway in the **Default gateway** field. Otherwise, leave the field blank.

9. In the **DNS Server** section, specify the DNS server information:

- If you added the PowerProtect Data Manager server hostname and IP address to the hosts file, then leave the default values in the **DNS Server** section.
- If the DNS server name is assigned automatically, select **Obtain DNS server address automatically**.
- If the DNS server IP address is static, select **Use the following DNS server addresses**, and then specify the IP address of the DNS server and any alternate DNS server that exists.

10. Verify the disk configuration, and then click **Next**.

NOTE: The disk size and number of hard disks that are added to the target machine should be either equal to or greater than that of the source machine.

11. Add the PowerProtect Data Manager server details, and then click **Next**.

NOTE:

- In **Server Name or IP**, enter only the FQDN or IP of the server.
- It is recommended that you use the Administrator role while providing PowerProtect Data Manager credentials.
- When a user with both the Restore Administrator and User roles runs the restore,
 - On the **Summary** tab, a prompt to enter Data Domain (DD) credentials appears. Enter the DD credentials and proceed with the restore process.
 - On the **Results** tab, the **Failed to unregister the host** error message appears. Click **OK** and verify the status of volumes.

If you provide a username with the Restore Administrator role and the restore fails with the message *Unable to fetch parameters required for performing restore*, see *File System Best Practices and Troubleshooting*.

12. On the **Select Backup** page, select the BMR data to restore to the host. Backups appear in the list in descending order from the most recent to the oldest.

13. Click **Next**.

A message is displayed while the information to complete the BMR operation is retrieved. Wait while the information is retrieved.

14. To add custom BMR options, click **Options** next to **Custom restore options**.

15. Perform one of the following actions:

CAUTION: If the **Restore physical machine to virtual machine (P2V)** option appears during the BMR wizard, select it. If you do not select it, the BMR operation succeeds, but the target host boots with a blue screen. For more information, see *Recover from a blue screen on boot after a BMR restore*.

- To accept the default PowerProtect Data Manager BMR options, click **Restore**.
- To specify non-default PowerProtect Data Manager BMR options, which are generally used for troubleshooting with assistance from Customer Support, perform these actions:

- a. Click **Options**.
- b. In the **Additional Options** field, type the options and values.

NOTE:

- o By default, for BMR restore the noncritical disks are not formatted. However, if you need to format any noncritical disks, you can use `-h DFA_SI_EXCLUDE_NON_CRITICAL_DYNAMIC_DISK=FALSE`. The default flag value of the `DFA_SI_EXCLUDE_NON_CRITICAL_DYNAMIC_DISK` additional option is set to **TRUE**.
- o Additional options must follow these guidelines:
 - Include a space between command and switch.
 - A key value pair should not contain any white space, for example,
`-h DFA_SI=TRUE -h DFA_SI_DR_P2V=TRUE`.

- c. Click **OK**.
- d. To confirm that you want to format the data and continue the BMR operation, click **OK**.
- e. Perform one of the following actions:
 - o To cancel the BMR operation, click **Cancel**.
 - o To proceed with the BMR operation, click **Restore**.

16. If you clicked **Restore**, wait until the BMR operation is successful.

NOTE: The wizard takes approximately 30 seconds to update the status (Cancelled, Failed, or Successful) in PowerProtect Data Manager.

To monitor the status of the BMR job from the PowerProtect Data Manager UI, select **Jobs > Protection Jobs**.

Next steps

Consider the following after the BMR operation is complete:

- The final status of the BMR operation is displayed on the **Results** tab.
- To open the BMR logs, click **View Logs**.
- To open a specific BMR log, select it, and then click **Open**.
- If the status of the recovery is **Cancelled** or **Failed**, restart the target host or boot it with the custom WinPE image again.
- If the virtual machine that was recovered boots using UEFI, the following error message appears in the `ddfarc` log file:
Virtual Disk Service error: This disk is already online. This message can be ignored.

Saving bare-metal recovery logs

The BMR logs might be needed for troubleshooting purposes. However, the WinPE environment does not allow copy, paste, or remote desktop connections.

About this task

To save the logs after the BMR operation completes, perform the following steps:

CAUTION: If the target host is restarted, the logs are lost.

Steps

1. If not already open, open a command-line window.
2. Mount a shared drive location to which you intend to copy the logs by running the following command:
`net use s:\<share-ip-address\sharename> /user:<username> <password>`
3. Run `cd X:\Program Files\DPSAPPS\fsagent\logs`.
4. Copy the files or folder to the mounted shared drive by running the following command:
`copy ddfarc.log s:\<name of destination folder on shared drive>`

Results

NOTE: You can also run `notepad.exe` to open the logs and see them in the WinPE environment.

Recover from a blue screen on boot after BMR

If you are performing a physical-to-virtual (P2V) operation with the **PowerProtect Data Manager Bare Metal Recovery Wizard**, the **Restore physical host to virtual host (P2V)** option might appear. If you do not select it, BMR succeeds but the target host boots with a blue screen.

About this task

In rare circumstances, network connectivity issues can prevent the BMR wizard from detecting if a target host is virtual or physical. If no confirmation is given, certain registry entries are not modified to the values required for a virtual host.

To correct the situation, perform the following actions:

Steps

1. Boot the target host with the WinPE image.
2. From the command-line window, run `diskpart` and `list volumes` to identify the current drive letter for the original system drive.
3. Run `ddfsrc.exe -h DFA_SI_DR_PATCH_REG_PATH="C:\Windows"`, replacing C: with the drive letter obtained in step 2.
4. Restart the host.

Next steps

The target host boots without a blue screen.

Perform bare-metal recovery of Windows clusters

Bare-metal recovery (BMR) can restore a Windows cluster configuration. BMR restores only the system state and critical disks on the cluster nodes.

Prerequisites

- The shared disk data and user data are protected using the relevant application agents.
- The hardware configuration of the cluster-node Windows clients is recorded. Necessary information includes the hardware vendor, size and type of disks, type of NIC, and amount of memory assigned.

About this task

To restore the system state and critical disks on the cluster nodes, perform the following steps.

Steps

1. If the old cluster-node Windows clients are powered on, power them off.
2. Create two new Windows clients with the same hardware configuration as the old cluster-node Windows clients.
3. Perform BMR on both of the new Windows clients.
4. Power on the Windows clients.
5. Open the Failover Cluster Manager and wait for the cluster to be connected.
6. Open Disk Management and initialize the new critical disks with the same configuration and labels as the old critical disks.
7. For each non-Cluster Shared Volume disk, perform the following substeps:
 - a. Perform a file-system restore.
 - b. Select the disk, and then click **More Actions > Repair**, replacing the disk.
 - c. Select the disk, and then click **Bring Online**.
8. Refresh the cluster node and confirm that all non-Cluster Shared Volume disks are online.
9. In the Failover Cluster Manager, perform the following substeps for each Cluster Shared Volume:
 - a. Record its name.
 - b. Remove it.
10. For each Cluster Shared Volume disk, perform the following substeps:
 - a. Select the disk, and then click **More Actions > Repair**, replacing the disk.

- b. Perform a file-system restore for each Cluster Shared Volume, using the name recorded in step 9.
 - c. Select the disk, and then click **Bring Online**.
11. Refresh the cluster node and confirm that all Cluster Shared Volume disks are online.

Next steps

Use normal File System agent backups to restore file system data, noncritical disks, and critical disks on shared storage. Also, use application-agent backups to restore application data.

Disaster recovery solution for Microsoft SQL application data

This section provides an end-to-end Disaster Recovery (DR) solution for Microsoft SQL application data in Microsoft Windows. This solution requires Microsoft SQL Server application and File System agents on Windows.

File System agent on Windows (FSA-BMR) protects Windows System data (critical volumes) and, the Microsoft SQL Server application agent protects the Microsoft SQL Server System and user databases. This scope of protection derives from Microsoft.

In detail, the volumes, where Microsoft SQL Server is installed and configured are considered as critical volumes. The FSA-BMR protects only the Windows critical volumes, which include Windows System Partition and any other volumes, where the third-party application or service is installed. Also, through FSA-BMR, the Windows System data, SQL installation, and configuration data are protected. The Microsoft SQL Server application agent protects the Microsoft SQL Server system database and user-created databases.

Protecting Windows system data and its critical volumes

Install and configure the File System agent on Windows to protect the Windows System data and critical volumes. See *Critical volumes in disaster recovery* for more information about the steps to configure and protect Windows system data.

Protecting Microsoft Server SQL application data

Install and configure the Microsoft SQL Server application agent to protect the Microsoft SQL Server system and user database. See *PowerProtect Data Manager Microsoft SQL Server User Guide* for more information about the steps to configure and protect the Microsoft SQL Server system and user database.

Performing DR for Microsoft SQL application data

The Microsoft SQL Server application agent needs a Microsoft SQL Server service to be running in single-user mode to restore the system database (master database). Since the File System agent on Windows does not backup the SQL server system database, the user needs to perform the below steps for an end-to-end DR solution for the SQL system database and user database in Windows.

About this task

See Microsoft documentation and *PowerProtect Data Manager Microsoft SQL Server User Guide* for more information on the following procedure.

Steps

1. Perform the BMR restore of Windows. See *Performing bare-metal recovery* for more information.
2. Mount the Microsoft SQL Server installation media in Microsoft Windows and run the following command. This rebuilds the Microsoft SQL Server System database.

```
Setup /QUIET /ACTION=REBUILDDATABASE /INSTANCENAME=MSSQLSERVER_2017 /  
SQLSYSADMINACCOUNTS=bmrtestdomain\administrator /SAPWD=Password
```

NOTE: Select the Microsoft SQL Server version according to the SQL version instance installed and configured. The above command was for the Microsoft SQL Server version 2017.

3. Start the Microsoft SQL Server instance in single-user mode and using the Microsoft SQL Server application agent, restore the Microsoft SQL server system database (master database). There are many ways, the user can start the Microsoft SQL Server instance in single-user mode. Perform the following steps to configure Microsoft SQL Server 2017 instance in a single-user mode.
 - a. Start the SQL Server Configuration Manager.
 - b. Right-click on the SQL server instance and select **Properties**.
 - c. In the **Startup parameters** tab, enter **-m** in the **Specify a startup parameter** field and click **Add**.
 - d. Click **Apply** and **Ok**.
 - e. Restart the SQL server instance.

The following example describes restoring the System DB through CLI:

```
C:\Program Files\DPSAPPS\MSAPPAGENT\bin>ddbmsqlrc.exe -a
NSR_DFA_SI_DD_HOST=192.162.1.1 -a NSR_DFA_SI_DD_USER=PLC_SQL_62-ppdm1461-flcad
-a NSR_DFA_SI_DEVICE_PATH=/PLC_SQL_62-ppdm1461-flcad/PLCTLIP-fddae07f-d1c3-497e-8b7c-
a8d90f270812 -a "NSR_DFA_SI_DD_LOCKBOX_PATH=C:\Program Files\DPSAPPS\common\lockbox"
-c bmrtestvm-62.bmrtestdomain.com -a "SKIP_CLIENT_RESOLUTION=TRUE" -f -t "06/03/2022
08:08:36 PM" -S normal MSSQL$MSSQLSERVER_2017:master
```

4. Start the Microsoft SQL Server instance in multiuser mode. Use the Microsoft SQL Server application agent to restore the msdb and model database.
 - a. Start the SQL Server Configuration Manager.
 - b. Right-click on the SQL server instance and select **Properties**.
 - c. In the **Startup parameters** tab, select **-m** from the **Existing parameters** field and click **Remove**.
 - d. Click **Apply** and **Ok**.
 - e. Restart the SQL server instance.

See *PowerProtect Data Manager Microsoft SQL Server User Guide* follow the steps to recover the Microsoft SQL Server msdb and model database.

The following example describes restoring the msdb database through CLI:

```
ddbmsqlrc.exe -a NSR_DFA_SI_DD_HOST=192.168.1.1 -a NSR_DFA_SI_DD_USER=PLC_SQL_62-
ppdm1461-flcad -a NSR_DFA_SI_DEVICE_PATH=/PLC_SQL_62-ppdm1461-flcad/PLCTLIP-fddae07f-
d1c3-497e-8b7c-a8d90f270812 -a "NSR_DFA_SI_DD_LOCKBOX_PATH=C:\Program
Files\DPSAPPS\common\lockbox" -c bmrtestvm-62.bmrtestdomain.com -a
"SKIP_CLIENT_RESOLUTION=TRUE" -f -t "06/03/2022 08:08:36 PM" -S normal
MSSQL$MSSQLSERVER_2017:msdb
```

The following example describes restoring the model database through CLI:

```
ddbmsqlrc.exe -a NSR_DFA_SI_DD_HOST=192.168.1.1 -a NSR_DFA_SI_DD_USER=PLC_SQL_62-
ppdm1461-flcad -a NSR_DFA_SI_DEVICE_PATH=/PLC_SQL_62-ppdm1461-flcad/PLCTLIP-fddae07f-
d1c3-497e-8b7c-a8d90f270812 -a "NSR_DFA_SI_DD_LOCKBOX_PATH=C:\Program
Files\DPSAPPS\common\lockbox" -c bmrtestvm-62.bmrtestdomain.com -a
"SKIP_CLIENT_RESOLUTION=TRUE" -f -t "06/03/2022 08:08:36 PM" -S normal
MSSQL$MSSQLSERVER_2017:model
```

5. Restore the user database using the Microsoft SQL Server application agent.
See *PowerProtect Data Manager Microsoft SQL Server User Guide* follow the steps to recover the Microsoft SQL Server user database.

Protecting Microsoft Distributed File System using BMR and SSR

Bare-metal recovery (BMR) and system-state recovery (SSR) are responsible for protecting the Microsoft Distributed File System (DFS) metadata/namespaces.

DFS data can be categorized into,

- DFS metadata (DFS namespace)
- DFS user data (user data share linked to DFS links)

Depending on the DFS configuration type (stand-alone or domain-based DFS), DFS stores its metadata in various locations such as registry, active directory, and so on.

BMR backup	If DFS links and DFS user data share are part of critical volumes, BMR backups them. See Critical volumes in disaster recovery for more information.
BMR restore	If DFS links and DFS user data are part of critical volumes, by default, BMR restores them including all the DFS metadata.
SSR backup	SSR backups only the DFS config data but not the DFS user data.
SSR restore	Always perform the full system state restore to fetch the DFS metadata information. If the DFS machine is a domain-based appliance, perform the full system state restore followed by AD authoritative restore.
	NOTE: Performing full system state restore helps in completing the restore process even when DFS service files are corrupted.

To protect DFS user data,

- Use File System agent to protect user data.
- If the user data is present on NAS share, use NAS agent to protect it.
- If the user data is part DFS other than File System or NAS agents, use the respective agent to protect it.

Performing application restores after bare-metal recovery

Some applications are not recovered by BMR.

When you back up BMR data, the backups include binaries for applications that use Windows services, such as Microsoft SQL. However, the backups normally exclude binaries for applications that do not use Windows services, as well as their configuration, databases, and file.

You need to reinstall the following applications and their data, or restore them from a File System agent backup:

- Applications that do not use Windows services.
- Applications installed to a noncritical volume that has been destroyed.

To restore application data, use the relevant application agent backup.

File System Best Practices and Troubleshooting

Topics:

- Troubleshooting installation and operation
- Troubleshooting backups
- Troubleshooting disaster recovery
- Troubleshooting restores
- Troubleshooting storage units

Troubleshooting installation and operation

You might encounter the following issues while installing or operating the File System agent.

PowerProtect agent service fails to start on Windows

After the application agent is installed, the PowerProtect agent service might fail to start on the Windows operating system. Check for the following timeout error messages in the event viewer logs:

```
A timeout was reached (30000 milliseconds) while waiting for the PowerProtect Agent Service service to connect.
The PowerProtect Agent Service service failed to start due to the following error:
The service did not respond to the start or control request in a timely fashion.
```

If these error messages appear in the event viewer logs, apply the following workaround:

<https://support.microsoft.com/en-in/help/922918/a-service-does-not-start-and-events-7000-and-7011-are-logged-in-window>

Agent registration

On Windows, if the agent fails to establish a connection with the PowerProtect Data Manager server, agent registration might fail with the following error message:

```
During a network connectivity test, the agent is unable to reach the PowerProtect Data Manager server by using ping.
```

1. If the ping command is blocked in the environment, the agent registration can still complete successfully. Review the agent service logs at `INSTALL_DIR\DPSAPPS\AgentService\logs` to verify that the registration is successful. If the registration is successful, the status of the agent host indicates **Registered** in the PowerProtect Data Manager UI.
2. If the ping command is not blocked in the environment, the agent registration might not complete successfully because a network connection cannot be started. If this occurs, complete the following steps to troubleshoot the issue:

On Linux or AIX, if the agent fails to establish a connection with the PowerProtect Data Manager server, agent registration might fail with the following error message:

```
During a network connectivity test, the agent is unable to reach the PowerProtect Data Manager server by using ping and curl.
```

1. If the ping command is blocked in the environment and curl is not installed, the

agent registration can still complete successfully.
Review the agent service logs at `/opt/dpsapps/agentsvc/logs` to verify that the registration is successful. If the registration is successful, the status of the agent host indicates **Registered** in the PowerProtect Data Manager UI.
2. If the ping command is not blocked in the environment, the agent registration might not complete successfully because a network connection cannot be started. If this occurs, complete the following steps to troubleshoot the issue:

If agent registration fails with these error messages, complete the following operation:

1. Use any network packet tracing tool to trace the packets from the agent system to PowerProtect Data Manager.
2. Start the packet tracing between the source IP of the agent system and the destination IP of PowerProtect Data Manager.
3. Start the network traffic between the agent system and PowerProtect Data Manager.

Wait 10 to 15 seconds.

4. Analyze the captured packets.
5. Look for SYN and SYN_ACK packets to see if a 3-way handshake is being performed.

Determine whether the source agent or the destination PowerProtect Data Manager is blocking the connection.

If network traffic is blocked, contact your network security team to resolve the port communication issue.

Agent host is registered with short name instead of FQDN

Application agent registration using the short name of the agent host is not supported after updating to PowerProtect Data Manager 19.11 or later versions.

If the agent host was previously registered with the short name, and you attempt to reregister the host with the short name after updating PowerProtect Data Manager to version 19.11 or a later version, registration fails with the following error:

```
Unable to register the agent host <hostName> because short name is not allowed for agent registration.
```

To resolve the issue, do not remove the host. Configure the host with the Fully Qualified Domain Name (FQDN), and then retry the registration.

If you try to register a new host with the short name after updating PowerProtect Data Manager to version 19.11 or a later version, registration fails with the following error:

```
Unable to register the agent host <hostName> because short name is not allowed for agent registration.
```

To resolve this issue, remove the application agent host in the PowerProtect Data Manager UI:

1. Go to **Infrastructure > Application Agents**.
2. Select the entry for the application agent host and click **Remove**.

After the removal, the host is in a "Deleted" state.

NOTE: The application agent host is automatically removed from the **Application Agents** window within 24 hours after the deletion.

3. After the removal process is complete, configure the host with the Fully Qualified Domain Name (FQDN), and then retry the registration.

PowerProtect agent service operations

To troubleshoot PowerProtect agent service operations, you can check the PowerProtect agent service log file `OpAgentSvc-<timestamp>.log`, which is created in `<agent_service_installation_location>\logs` on Windows and `<agent_service_installation_location>/logs` on AIX or Linux. To modify the log level and retention of temporary files, you can modify specific parameter settings in the `config.yml` file.

To modify the log level and retention of temporary files, you can perform the following steps:

1. Stop the agent service.
2. Open the `config.yml` file in an editor.
3. Modify the log-level settings in the following parameters, as required:
 - `DEBUG`

- INFO
- WARNING
- ERROR
- CRITICAL

NOTE: These parameters are listed in order of decreasing number of messages in the debug information output. The default log-level is INFO.

- To retain the temporary files, set the `keepTempFiles` parameter to `True` in the `config.yml` file.

NOTE: The agent service and application agent communicate through the temporary files, which are typically deleted after use but can be useful for troubleshooting purposes. Do not leave the `keepTempFiles` parameter set to `True` permanently, or the temporary files can use excessive space on the file system.

- Start the agent service.

Uninstalling the PowerProtect agent service results in "Segmentation fault(coredump)" error on AIX platforms

After updating the application agent from version 19.12 to 19.13, uninstalling the PowerProtect agent service results in the following error:

```
/usr/sbin/rpm_share [470]: 19005466 Segmentation fault(coredump)
```

This behavior does not impact reinstallation of the PowerProtect agent service.

The core dump error is due to a faulty `rpm.rte` file with version 4.13.0.3. For more information, see update `rpm.rte` to the latest version to resolve `rpm` errors and core dumps.

To resolve this issue, either update the `rpm.rte` file to the latest version or update the AIX Technology Level (TL) version. The following article provides more information about TL versions:

Fileset information for: `rpm.rte`

Deletion compliance for SLA validation job marked as "Failed" in PowerProtect Data Manager UI

Due to a cancel request and create-new-copy deletion request occurring simultaneously, deletion compliance for an SLA validation job is incorrectly shown as "Failed" in the UI, even though the copy deletion has succeeded. This issue can occur if replication copies exist, but the corresponding backup copy has been deleted. In this case, the retention time of the replication copies is greater than the retention time of the backup copy. Once the replication copies expire or are manually deleted, it triggers the catalog deletion request, which marks the copy as "Deleted."

PowerProtect Data Manager UI display of localhost.localdomain hostname

In the PowerProtect Data Manager UI, the **Application Agents**, **Asset Sources**, and **Protection Jobs** windows might list the asset primary hostname as `localhost.localdomain` instead of the expected FQDN.

The display of `localhost.localdomain` as the hostname in the PowerProtect Data Manager UI windows might occur when you specify the host's actual FQDN setting for the loopback address in the `/etc/hosts` file. For example, when you add the following settings in the `/etc/hosts` file, the first setting value, `localhost.localdomain`, appears as the hostname in the PowerProtect Data Manager UI windows, instead of the actual FQDN:

```
127.0.0.1 localhost.localdomain localhost6 localhost6.localdomain6
127.0.0.1 blrv027d233.blr.lab.dell.com blrv027d233
```

Ensure that the host's actual FQDN is not specified for the loopback address and do not specify hostnames that start with "local" in the `/etc/hosts` file.

Troubleshooting backups

You might encounter the following issues while performing backups with the File System agent.

Backing up Windows ACL properties

On a Windows client, if the contents of a file have not changed since the last backup but the Access Control List (ACL) properties have, incremental backups will not back up the changed ACL properties.

Perform the following steps to back up the changed ACL information:

1. Enable Last Access from the client by running the command `fsutil behavior set disablelastaccess 0`.
2. Restart the client.
3. Check the status by running `fsutil behavior query disablelastaccess` and looking for `DisableLastAccess = 0` in the output.
4. Set the `detect-acl-changes` flag to `true`. The value of this flag is `false` by default. Edit the `C:\Program Files\DPSFSAGENT\settings\ddfsav.fsagentconfig` file and change "`--detect-acl-changes=true`" to set the flag.

Take note of the following:

- When the `--detect-acl-changes` flag is set, the file is treated as a modified file and is backed up as part of the next backup.
- If ACL modifications are made only to executable files residing on a mounted volume, the changes might not be backed up.
- This works only for files, not folders. Changes only to the ACL properties of folders cannot be backed up.

Backups fail when credentials include a backslash character (\)

When you enter credentials that include a backslash character (\) for an application agent in the PowerProtect Data Manager UI, the backups fail.

For example, when you enter a password for the operating system or database user that includes the backslash character, subsequent backups fail with the following error message:

```
systemErr: Unable to log in.
```

This error might occur when updating the password for a storage unit.

To resolve this issue, type `\\` (double backslash) instead of `\` (single backslash) when you enter the credentials for an application agent in the PowerProtect Data Manager user interface.

Block-based backup driver installation

The following message might appear during the installation of the block-based backup driver: `Block based backup driver was installed but not loaded`. If the driver cannot be loaded, file-based backups will be performed instead of block-based backups.

Perform the following troubleshooting steps in order. Unless otherwise noted, if you make a change at any step, test block-based backups again before proceeding to the next step:

1. If the `bc` utility is not installed, install the utility.
2. If the `livepatch` module is installed and enabled, disable it.
3. Install or reinstall the block-based driver `ppdm-bbbwt`.
4. Restart the `nsrbbb` service. For information on the specific commands to execute, refer to the documentation of the host's operating system.

NOTE: If you are unable to resolve the issue, contact Customer Support.

Linux block-based incremental backups fail

On Linux, the block-based incremental backups consistently fail and display a message similar to `save: Block Based Error subsystem error while performing Block Based Backup.`

Check if any other process is already accessing the snapshot or delete the snapshot manually, and then try again.

Edit Retention button is disabled for asset copies

On the **Infrastructure > Assets page > Asset <asset name>** page, the **Edit Retention** button is disabled.

The **Edit Retention** button is enabled only when the copy status of the asset source is **Available**. In detail, if deleting a copy of a protected asset fails, the **Copy Status** will be updated to another state from **Available**, and the **Edit Retention** button is disabled.

Workarounds are:

- Edit the retention period of an asset copy if its status is **Available**.
- Do not try deleting an asset copy before its retention period.

Troubleshooting disaster recovery

You might encounter the following issues while performing disaster recovery with the File System agent.

BMR and SSR operations fail to run and a `start_subscriber` error appears in the agent-service log

A BMR or SSR operation might fail to display the start or completion of a job, and the `DPSAPPS\AgetService` log shows an error similar to `msg_server.py-create_subscriber()Line 617 Exception occurred during start_subscriber Permission denied occurred.`

This is caused when another service uses port 7010 or 7011.

Reconfigure any service that uses port 7010 or 7011 to use a different port.

Backing up BMR data fails with a `ddfssv` error

Backing up BMR data can fail with the error `ddfssv FATAL <12128: Attempt to create a backup with no data is unsuccessful.`

The Event Viewer system log might also contain an error similar to the following:

The shadow copies of volume H: were deleted because the shadow copy storage could not grow in time. Consider reducing the IO load on the system or choose a shadow copy storage volume that is not being shadow copied.

This is a known Microsoft issue that affects all backup products.

Retry the backup procedure after leaving more free space on the drive or when the drive is not as busy.

Disaster-recovery backups fail with a `The system cannot find the file specified` error

Disaster-recovery backups can fail with error messages similar to the following in the `ddfssv` log file:

```
2022-09-08T14:00:04.687Z
[TRACE_ID:8815fbc662dd39f9;JOB_ID:af8c2587f740e6c4;EXEC_ID:a869d8aa1b7aaf57] ddfssv
Error <5737>: I/O error: Unable to BackupRead(3) "C:\System Volume
Information\EfaSIDat\SYMEFA.DB" (code 2: The system cannot find the file specified).
```

```
2022-09-08T14:00:04.687Z
[TRACE ID:8815fbc662dd39f9;JOB_ID:af8c2587f740e6c4;EXEC_ID:a869d8aalb7aaf57] ddfssv
Error <16507>: Path "C:\System Volume Information\EfaSIDat\SYMEFA.DB": Total processed
bytes 0 is inconsistent with original data size 18087936
2022-09-08T14:00:12.499Z
[TRACE ID:8815fbc662dd39f9;JOB_ID:af8c2587f740e6c4;EXEC_ID:a869d8aalb7aaf57] ddfssv
Error <5137>: Unable to open 'C:\System Volume Information\EfaSIDat\SYMEFA_1.DB' for
backup: code 5: Access is denied.
```

This can be caused by antivirus software preventing access to system files.

To resolve this issue, perform one of the following tasks:

- Use file exclusion:
 1. Exclude `ddfssv.exe`, `ddfsrc.exe`, `ddfscpn.exe`, and `restserver.exe` from the antivirus software.
 2. Retry the disaster-recovery backup.
 3. Optionally, remove the files from the exclusion list.
- Temporarily disable the antivirus software:
 1. Disconnect the host from the network.
 2. Disable the antivirus software.
 3. Retry the disaster-recovery backup.
 4. Enable the antivirus software.
 5. Connect the host to the network.

Network connectivity issues after a BMR

A BMR can fail to recover network configuration information such as IP address, subnet mask, gateway, and DNS. If the recovered host experiences network connectivity issues, confirm its network configuration and make any necessary changes.

SSRs fail to restore a VSS writer

If an SSR fails to recover a VSS writer, you might see an error similar to the following:

```
ddfsrc Error <0000>: Unable to select W component from V for restore: The specified
object was not found. (VSS error 0x80042398)
```

This occurs when writer .xml files are missing from the `C:\Windows\Vss\Writers\System` directory.

Restart the host, run the command `vssadmin list writers` to confirm the VSS writer status, and then retry the SSR.

BMR recovers critical volumes, but a disk is marked as offline

This can happen with disks that host Microsoft SQL Server instances.

Manually bring the disk online, and then restart the host.

SSRs show a `RegSetValueEx()` error and fail to recover files after a restart

An SSR might show `Error <10958>: RegSetValueEx() for replace files`, but complete successfully. After the host is restarted, certain files have not been recovered.

A file that is in use cannot be replaced by its recovered version. If you restart a host to have a file replaced but the file is not replaced, it might be because antivirus software is preventing access to the `pending/renamoperations` registry key.

Perform one of the following tasks as a workaround.

- Use file exclusion:
 1. Retry the SSR, but do not restart the host.
 2. Exclude `ddfssv.exe`, `ddfsrc.exe`, `ddfscpn.exe`, and `restserver.exe` from the antivirus software.
 3. Restart the host.

4. Optionally, remove the files from the exclusion list.
- Temporarily disable the antivirus software:
 1. Retry the SSR, but do not restart the host.
 2. Disconnect the host from the network.
 3. Disable the antivirus software.
 4. Restart the host.
 5. Enable the antivirus software.
 6. Connect the host to the network.

Temporary files can be deleted manually after a recovery operation

The temporary files that are generated during the recovery operation are stored in `C:\dbapps_temp_dir`. You can delete these temporary files manually after the recovery operation and the reboot.

NOTE:

- Replace `C:` with the drive to which Windows is installed.
- You may require special permissions to delete some temporary files. To delete such files, contact Customer Support.

BMR restore errors

When performing a BMR restore, you might see the errors similar to the following:

```
Assigning volume letters or mount points failed.
```

```
Unable to do VSS restore: Bare metal recovery remount volumes step failed.
```

Perform the BMR restore again.

Cluster assets backup failure

The cluster assets backup fails when there is a combination of either stand-alone and cluster assets, or disaster recovery and cluster assets in the same policy.

Workaround is to run the cluster backup in a different policy and schedule.

BMR backup fails on EFI based system for a few times when backing up the BCD files

Backup for Extensible Firmware Interface (EFI) partition is performed from the live volume as the snapshot volume of the EFI partition cannot be created. If the system process uses the Boot Configuration Data (BCD) files, the `ddfssv` process tries to export a temporary copy for BCD files using the Microsoft `bcdutil` tool. This issue occurs if the export operation causes any change in the exported file.

A sample `ddfssv` log file with errors:

```
-----  
2022-05-18T08:10:08.914Z  
[TRACE ID:a559c76a1e886397;JOB_ID:9333c654df01e5af;EXEC_ID:af0c561e68d3aa75] ddfssv Info  
<0000>: Exported bcd file to temporary location  
2022-05-18T08:10:08.920Z  
[TRACE ID:a559c76a1e886397;JOB_ID:9333c654df01e5af;EXEC_ID:af0c561e68d3aa75] ddfssv  
Error <16507>: Path \\?\GLOBALROOT\Device\HarddiskVolume2\EFI\Microsoft\Boot\BCD: Total  
processed bytes 32768 is inconsistent with original data size xxxxx  
2022-05-18T08:10:08.926Z  
[TRACE ID:a559c76a1e886397;JOB_ID:9333c654df01e5af;EXEC_ID:af0c561e68d3aa75] ddfssv  
Error <16507>: Path \\?\GLOBALROOT\Device\HarddiskVolume2\EFI\Microsoft\Boot\BCD.LOG:  
Total processed bytes 32768 is inconsistent with original data size xxxxx  
...
```

```
2022-05-18T10:51:25.0782 [TRACE_ID_a559c76ale886397_EXEC_ID_af0c561e68d3aa75] ddfasv
Info <6149>: Error summary: 2 errors: 16507(2)
2022-05-18T10:51:25.1022 [TRACE_ID_a559c76ale886397_EXEC_ID_af0c561e68d3aa75] ddfasv
Info <5314>: Command completed with exceptions (2 errors, exit code 10020: completed
with errors, client log should be examined)
...
-----
```

Workaround is to rerun the backup.

IPv6 configuration fails in WinPE

When the user modifies the IPv6 address during the host configuration in the WinPE wizard, the system adds and lists all the IPv6 addresses rather than replacing the existing IPv6 address.

Workaround is to restart the restore process.

BMR fails with the message Unable to fetch parameters required for performing restore

If the BMR wizard fails with this message, and you have provided a username with the Restore Administrator role, ensure that the user account also has the User role.

Backup host is unregistered in PowerProtect Data Manager when the same FQDN and IP address are used in BMR

If the fully qualified domain name (FQDN) and IP address used in registering the backup host in PowerProtect Data Manager are used during bare-metal recovery (BMR) in the WinPE wizard, the registered backup host is unregistered after the BMR restore. Also, the same host is registered as a new one in PowerProtect Data Manager.

NOTE: Ensure that you use a different FQDN and IP address in the WinPE wizard during BMR to prevent the backup host from unregistering from PowerProtect Data Manager.

As a workaround, perform the following:

1. Unregister the new host from PowerProtect Data Manager using the unregister.bat file.
2. Reregister the backup host in PowerProtect Data Manager using the register.bat file.

NOTE:

- Both register.bat and unregister.bat files are available in C:\Program Files\DPSAPPS\AgentService.
- If required, after reregistering the backup host, manually discover and add assets to the existing protection policy.

Troubleshooting restores

You might encounter the following issues while performing restores with the File System agent.

File-level restores of symbolic-linked directories

The file-level restore of a symbolic-linked directory does not contain any symbolic links or reparse points, but the contents are the same as the source directory.

Create symbolic links manually after the file-level restore operation completes.

File-level restore operations fail with a client services error

When performing a file-level restore, you might see an error similar to one of the following:

```
Resource not found. Check if client services are running and perform the operation again.
```

```
ARA0002: Unable to restore FILE_SYSTEM asset F:\ because of an agent issue.
```

```
The restore was unsuccessful because of an agent issue.
```

```
The File System restore was unsuccessful because a Status code 500 error occurred during the request to the Agent host.
```

Perform the operation again.

File sparseness and Linux file-level restores

Linux file-level restores of file-level backups do not retain file sparseness.

Use the appropriate tools to reduce disk space.

Restores of clustered drives fail with an agent host timed out error or a network connectivity error

When attempting to restore a clustered drive, you might see an error similar to one of the following:

```
The File System restore was unsuccessful because the request to the Agent host timed out.
```

```
ARA0005: Unable to restore FILE_SYSTEM asset C:\ClusterStorage\Volume3 because of a network connectivity issue on agent host w2019c2.agent.com.
```

```
The restore was unsuccessful because of an issue with the network connection.
```

```
To resolve this issue: 1.Check the network connectivity between PowerProtect Data Manager and the agent host. 2.Confirm that there is no packet loss between PowerProtect Data Manager and the agent host.
```

Perform a full discovery of the logical cluster host, and then retry the restore. For more information about performing a full discovery, see the *PowerProtect Data Manager Administrator Guide*.

Restoring block-based backups on 16 TB ReFS volumes

Image-level restore operations of File System agent block-based backups (BBB) might fail ReFS volumes with 16 TB or higher capacity.

Restore the data by using file-level restore (FLR).

Restoring multiple backups or SSIDs from the command-line interface

A single-level restore command-line interface (CLI) cannot be used to select multiple backups or SSIDs for restore.

Open a separate file-level restore CLI for each backup or SSID.

XFS restores of block-based backups

When performing an XFS restore of a block-based backup on a remote host, the kernel version of the remote host must be equal to or higher than the kernel version of the source host.

There is no workaround.

Troubleshooting storage units

Review the following issues related to storage units in PowerProtect Data Manager.

Creating storage unit fails when maximum MTree and Users count on DD system reached

When you add a protection policy or create a storage unit in PowerProtect Data Manager, storage unit creation fails if you reach the maximum MTree and Users count on the selected DD system. PowerProtect Data Manager enables you to finish adding a protection policy without a storage unit. However, if you subsequently run a backup with this protection policy, the backup process is suspended indefinitely with no error message.

To continue backup operations, you must perform a cleanup on the DD system.

Discrepancy between storage unit capacity reported in PowerProtect Data Manager and DD Virtual Edition

Due to differences in space calculation (physical capacity vs. logical capacity), there is a discrepancy between storage unit capacity reported in PowerProtect Data Manager and DD Virtual Edition. For example, the DD storage unit capacity displayed in the **Protection > Storage > Manage Storage** window of the PowerProtect Data Manager UI might be greater than the amount displayed in DDVE.

To determine storage unit capacity, use DDVE instead.

Glossary

This glossary provides definitions of acronyms used across the product documentation set.

A

AAG: Always On availability group
ACL: access control list
AD: Active Directory
AKS: Azure Kubernetes Service
API: application programming interface
ARM: Azure Resource Manager
AVS: Azure VMware Solution
AWS: Amazon Web Services
AZ: availability zone

B

BBB: block-based backup

C

CA: certificate authority
CBT: Changed Block Tracking
CDC: change data capture
CIFS: Common Internet File System
CLI: command-line interface
CLR: Common Language Runtime
CN: common name
CPU: central processing unit
CR: custom resource
CRD: custom resource definition
CSI: container storage interface
CSV: Cluster Shared Volume

D

DAG: database availability group
DBA: database administrator

DBID: database identifier

DDMC: DD Management Center

DDOS: DD Operating System

DDVE: DD Virtual Edition

deploy

At Dell, virtual machines are deployed to virtual environments, while software components and hardware devices are installed. Both PowerProtect Data Manager and DDVE are virtual machines that are deployed. If you are searching this software guide for instances of `install` and not finding anything appropriate, search for `deploy` instead.

DFC: DD Boost over Fibre Channel

DNS: Domain Name System

DPC: Data Protection Central

DR: disaster recovery

DRS: Distributed Resource Scheduler

DSA: Dell security advisory

E

EBS: Elastic Block Store

EC2: Elastic Compute Cloud

eCDM: Enterprise Copy Data Management

ECS: Elastic Cloud Storage

EFI: Extensible Firmware Interface

EKS: Elastic Kubernetes Service

ENI: Elastic Network Interface

EULA: end-user license agreement

F

FC: Fibre Channel

FCD: first class disk

FCI: failover cluster instance

FETB: front-end protected capacity by terabyte

FLR: file-level restore

FQDN: fully qualified domain name

FTP: File Transfer Protocol

G

GB: gigabyte

At Dell, this is 2^{30} bytes.

Gb/s: gigabits per second

At Dell, this is 2^{30} bits per second.

GCP: Google Cloud Platform

GCVE: Google Cloud Virtual Edition

GID: group identifier

GLR: granular-level restore

GUI: graphical user interface

GUID: globally unique identifier

H

HA: High Availability

HANA: high-performance analytic appliance

HTML: Hypertext Markup Language

HTTP: Hypertext Transfer Protocol

HTTPS: Hypertext Transfer Protocol Secure

I

IAM: identity and access management

IDE: Integrated Device Electronics

IP: Internet Protocol

IPv4: Internet Protocol version 4

IPv6: Internet Protocol version 6

K

KB: kilobyte

At Dell, this is 2^{10} bytes.

L

LAC: License Authorization Code

LAN: local area network

M

MB: megabyte

At Dell, this is 2^{20} bytes.

ms: millisecond

MTU: maximum transmission unit

N

NAS: network-attached storage

NBD: network block device

NBDSSL: network block device over SSL

NDMP: Network Data Management Protocol

NFC: Network File Copy

NFS: Network File System

NIC: network interface card

NTFS: New Technology File System

NTP: Network Time Protocol

O

OS: operating system

OSS: open-source software

OVA: Open Virtualization Appliance

P

PCS: Protection Copy Set

PDF: Portable Document Format

PEM: Privacy-enhanced Electronic Mail

PIN: personal identification number

PIT: point in time

PKCS: Public Key Cryptography Standards

PSC: Platform Service Controller

PVC (cloud computing): private virtual cloud

PVC (Kubernetes): Persistent Volume Claim

R

RAC: Real Application Cluster

RAM: random-access memory

RBAC: role-based access control

ReFS: Resilient File System
REST API: representational-state transfer API
RHEL: RedHat Enterprise Linux
RMAN: Recovery Manager
RPO: recovery-point objective
RSA: Rivest-Shamir-Adleman

S

S3: Simple Storage Services
SaaS: software as a service
SAP: System Analysis Program Development
From the SAP website (2022), "the name is an initialism of the company's original German name: Systemanalyse Programmentwicklung, which translates to System Analysis Program Development. Today the company's legal corporate name is SAP SE - SE stands for *societas Europaea*, a public company registered in accordance with the European Union corporate law."
SCSI: Small Computer System Interface
SDDC: software-defined data center
SELinux: Security-Enhanced Linux
SFTP: Secure File Transfer Protocol
SLA: service-level agreement
SLES: SuSE Linux Enterprise Server
SLO: service-level objective
SPBM: Storage Policy Based Management
SQL: Structured Query Language
SRS: Secure Remote Services
SSD: solid-state drive
SSH: Secure Shell
SSL: Secure Sockets Layer
SSMS: SQL Server Management Studio
SSVs: System Stable Values

T

TB: terabyte
At Dell, this is 2^{40} bytes.
TCP: Transmission Control Protocol
TDE: Transparent Data Encryption

TLS: Transport Layer Security
TPM: Trusted Platform Module
TSDM: Transparent Snapshots Data Mover
T-SQL: Transact-SQL

U

UAC: user account control
UDP: User Datagram Protocol
UI: user interface
UID: user identifier

update

At Dell, software is updated and hardware is upgraded. If you are searching this software guide for instances of upgrade and not finding any, search for update instead.

UTC: Coordinated Universal Time

From Wikipedia (2022), "this abbreviation comes as a result of the International Telecommunication Union and the International Astronomical Union wanting to use the same abbreviation in all languages. English speakers originally proposed CUT (for 'coordinated universal time'), while French speakers proposed TUC (for 'temps universel coordonné')."

V

VADP: VMware vStorage API for Data Protection
VBS: virtualization-based security
VCF: VMware Cloud Foundation
vCLS: vSphere Cluster Service
vCSA: vCenter Server Appliance
VCSA: vCenter Server Appliance
VDI: Virtual Device Interface
vDisk: virtual disk
vDS: virtual distributed switch
vFRC: Virtual Flash Read Cache
VGT: Virtual Guest Tagging
VIB: vSphere Installation Bundle
VLAN: virtual LAN
VM: virtual machine
VMC: VMware Cloud
VMDK: virtual machine disk
VNet: virtual network

VPC: virtual private cloud
vRSLCM: vRealize Suite Lifecycle Manager
VST: Virtual Switch Tagging
vTPM: Virtual Trusted Platform Module
VVD: VMware Validated Design
vVol: virtual volume

W

WAN: wide area network

RecoverPoint for Virtual Machines 5.3.3

vSphere HTML5 Plugin Administrator's Guide

Notes, cautions, and warnings

① **NOTE:** A NOTE indicates important information that helps you make better use of your product.

⚠ **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

⚠ **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Contents

Figures	6
Tables	7
Preface	8
Chapter 1: Before you begin	10
Create your license files.....	10
Open RecoverPoint for VMs.....	11
Add license.....	12
Register for Customer Support.....	13
Chapter 2: Protecting VMs	14
Protect a VM in a new group.....	14
Protect a VM in an existing group.....	17
Protect multiple VMs in a new group.....	19
Protect multiple VMs in an existing group.....	21
Add a copy to an existing group.....	23
Chapter 3: Monitoring protection	25
The RecoverPoint for VMs Dashboard.....	25
Monitor system alerts.....	29
Monitor system events.....	30
Monitor system limits.....	34
Monitor system components.....	37
Monitor group and copy protection.....	38
Chapter 4: VM automation and orchestration	39
Create a bookmark.....	39
Automation.....	40
Automatic protection of newly added VMDKs.....	41
Provisioning copy VMDKs.....	41
Excluding a VMDK from replication.....	41
Automatic replication of VM hardware changes.....	42
Orchestration.....	42
VM start-up sequence.....	42
Group start-up sequence.....	45
Create a group set.....	46
Re-IP rules.....	48
Failover networks.....	50
Chapter 5: Recovering VMs	52
Test a copy.....	52
Failover to a copy.....	56

Recover production from a copy.....	62
Chapter 6: Monitoring recovery activities.....	67
Appendix A: Managing RecoverPoint for VMs.....	68
Managing your RecoverPoint for VMs licenses.....	68
Managing the plugin server.....	68
Changing the plugin server certificate.....	69
Changing a registered plugin server certificate.....	69
Managing vCenter Server registration with plugin server.....	70
Managing linked vCenter Server registration.....	71
Displaying the plugin server info.....	75
Collecting plugin server logs.....	75
Upgrading the plugin server.....	76
Managing vRPA clusters.....	77
Managing system component registration.....	79
Register vCenter Server to vRPA cluster.....	80
Managing ESX cluster registration.....	81
Managing journal datastore registration.....	82
Managing external host registration.....	83
Managing protected VMs.....	83
Managing VM protection policies.....	83
Stop protecting a VM.....	84
Managing consistency groups.....	85
Managing group protection policies.....	85
Managing group sets.....	89
Managing recovery activities.....	90
Manage testing.....	91
Manage failover and failback.....	95
Manage production recovery.....	98
Configuring email alerts and reports.....	100
Appendix B: Troubleshooting.....	102
Finding the vRPA cluster management IP.....	102
Collecting logs.....	102
Collecting logs from vRPA clusters.....	102
Collect plugin server logs.....	105
Collecting RecoverPoint for VMs splitter logs.....	105
Adding new VMDKs.....	105
Removing a VMDK.....	105
Automatically expanding copy VMDKs.....	106
Recovering from a cluster disaster.....	106
RecoverPoint for VMs licensing.....	107
Register RecoverPoint by email or phone.....	108
Creating VMkernel ports.....	108
Load balancing.....	109
Copy VM network configuration guidelines.....	110
Changing the network adapter configuration of a protected VM.....	112

Appendix C: Shared disks	113
Add VM to a host cluster.....	113
Add shared disk to VMs that belong to a consistency group.....	113
Recover production after deletion of production VM or shared disk.....	114
Failover after deletion of production VM or shared disk.....	114
Remove a VM from its host cluster	115
Unprotect a shared disk.....	115

Figures

1	The RecoverPoint for VMs Dashboard.....	25
2	RecoverPoint for VMs event logs.....	31
3	System limits.....	35
4	Consistency group limits.....	35
5	vRPA cluster limits.....	36
6	Splitter limits.....	36
7	Monitoring recovery activities.....	67
8	Plugin server vCenter Servers screen.....	70
9	Local linked vCenter example.....	72
10	Remote linked vCenter example.....	74

Tables

1	Virtual machine network settings available through the GUI.....	110
2	Network settings only available through the JSON file.....	112

As part of an effort to improve product lines, we periodically release revisions of software. Therefore, some functions described in this document might not be supported by all versions of the software currently in use. The product release notes provide the most up-to-date information on product features.

Contact your technical support professional if a product does not function properly or does not function as described in this document.

NOTE: This document was accurate at publication time. Go to Online Support (<https://www.dell.com/support>) to ensure that you are using the latest version of this document.

Purpose

This document includes conceptual information on managing a RecoverPoint for Virtual Machines system.

Audience

This document is intended for use by vSphere administrators who are responsible for managing the RecoverPoint for VMs system.

Related documentation

The following publications provide additional information:

- *RecoverPoint for Virtual Machines Release Notes*
- *RecoverPoint for Virtual Machines Installation and Deployment Guide*
- *RecoverPoint for Virtual Machines Flex Plugin Administrator's Guide*
- *RecoverPoint for Virtual Machines Deployment REST API Programming Guide*
- *RecoverPoint for Virtual Machines REST API Programmer's Guide*
- *RecoverPoint for Virtual Machines Security Configuration Guide*
- *RecoverPoint for Virtual Machines Scale and Performance Guide*
- *RecoverPoint for Virtual Machines CLI Reference Guide*
- *RecoverPoint for Virtual Machines RESTful API* at <https://developer.dell.com/apis>.

In addition to the core documents, we also provide white papers, technical notes, and demos.

Typographical conventions

This document uses the following style conventions:

Bold	Used for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)
<i>Italic</i>	Used for full titles of publications referenced in text
Monospace	Used for: <ul style="list-style-type: none">• System code• System output, such as an error message or script• Pathnames, filenames, prompts, and syntax• Commands and options

<i>Monospace italic</i>	Used for variables
Monospace bold	Used for user input
[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections - the bar means "or"
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Product documentation

- For release notes and user guides, go to **Online Support** at <https://www.dell.com/support>.
- For API documentation, see <https://developer.dell.com/apis>.

Product information

For documentation, release notes, software updates, or information about products, go to **Online Support** at <https://www.dell.com/support>.

Where to get help

Go to **Online Support** at <https://www.dell.com/support> and click **Contact Support**. To open a service request, you must have a valid support agreement. Contact your sales representative for details about obtaining a valid support agreement or with questions about your account.

Where to find the support matrix

Consult the **Simple Support Matrix** for RecoverPoint for VMs at <https://elabnavigator.emc.com>.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to <https://contentfeedback.dell.com/s>.

Before you begin

Before you start protecting your data in RecoverPoint for VMs, perform the tasks in this section in the provided sequence.

This guide provides the procedures for protecting, recovering and managing VMs:

- **Using the RecoverPoint for VMs Plugin for vSphere Client (HTML5) version 6.7 U1 or later.** When using the RecoverPoint for VMs Flex Plugin, see the *RecoverPoint for VMs Flex Plugin Administrator's Guide*.

Before you begin:

- See the *RecoverPoint for VMs Product Guide* for a detailed description of the **vSphere HTML5 plugin** and the **vSphere HTML5 plugin server**.
- See the *RecoverPoint for Virtual Machines Scale and Performance Guide* and the *RecoverPoint for Virtual Machines Release Notes* for information of how to scale your environment, and the limitations of this solution.
- System installation must be complete. See the *RecoverPoint for VMs Installation and Deployment Guide* for more information.
- The plugin server must be registered with the vCenter Server that you are connected to, or one linked to the vCenter you are connected to. See *Managing the plugin server* for more information.

Topics:

- Create your license files
- Open RecoverPoint for VMs
- Add license
- Register for Customer Support

Create your license files

When a RecoverPoint for VMs sales order is approved, an order confirmation email is automatically sent to the email addresses provided during order entry. The email provides the information you need to begin license activation.

About this task

For more information about software licensing, see these resources:

- Software licensing documentation
- Software Licensing Central Activation video

Steps

1. If you are starting from the Dell Digital Locker, log in or create an account, search for your order, select your product, and click **Activate now**.
This action takes you to the first step of the license activation wizard in the Software Licensing Center (described in step 3).
2. If you do not have the order confirmation email, you can activate your software directly from the Software Licensing Center:
 - a. Click **ACTIVATE MY SOFTWARE**.
 - b. Provide information such as Licensing Authorization Code (LAC) or Sales Order #.
3. In the **SELECT PRODUCTS** step of the activation wizard, select the product you want to activate, and click **START THE ACTIVATION PROCESS**.
4. In the **SELECT A MACHINE** step:
 - a. If you are activating a new instance or rehosting to a new target machine, add a new machine name, and click **SAVE MACHINE & CONTINUE TO NEXT STEP**.
 - b. If you are increasing capacity on an existing machine, use the **SEARCH MACHINES** widget, select the wanted machine, and click **NEXT: ENTER DETAILS**.
5. In the **ENTER DETAILS** step, enter the quantity of entitlements that you want to activate and the vCenter ID. To learn how to obtain the vCenter ID, click **Machine Details FAQs** and select **RecoverPoint for Virtual Machines**.

6. In the **REVIEW** step, review your selections, and click **ACTIVATE**. This action generates the license keys.
7. In the **COMPLETE** step, view and download the license keys that you generated. To return to the Software Licensing Center home page, click **HOME**.

Results

The entitlements are converted to license files.

Next steps

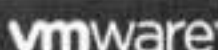
Transfer the license files to the computer from which you will be running RecoverPoint for VMs.

Open RecoverPoint for VMs

Display the RecoverPoint for VMs plugin in your vSphere Client.

Steps

1. Connect to a vCenter Server hosting RecoverPoint for VMs components.
2. Click **LAUNCH VSPHERE CLIENT (HTML5)**. You can also launch the **vSphere Client (HTML5)** directly by entering `https://vCenter-IP or FQDN/ui/` into your address bar. The HTML5 plugin supports vSphere 6.7 U1 and later versions. In vSphere 7.0 U2 or later, it is recommended to use RecoverPoint for VMs 5.3 SP2 or later.



Getting Started



Documentation

[VMware vSphere Documentation Center](#)

Best practice is to **LAUNCH VSPHERE CLIENT (HTML5)** as your primary client, and only use the **vSphere FLEX plugin** if you are using a vSphere version prior to 6.5 U1, or if you need a feature that is not currently supported through the **vSphere HTML5 plugin** of your RecoverPoint for VMs version. See the *RecoverPoint for VMs Flex Plugin Administrator's Guide* for more information.

In RecoverPoint for VMs 5.3 SP2 and later versions, when running a vSphere version that supports FLEX, only use the **vSphere FLEX plugin** to:

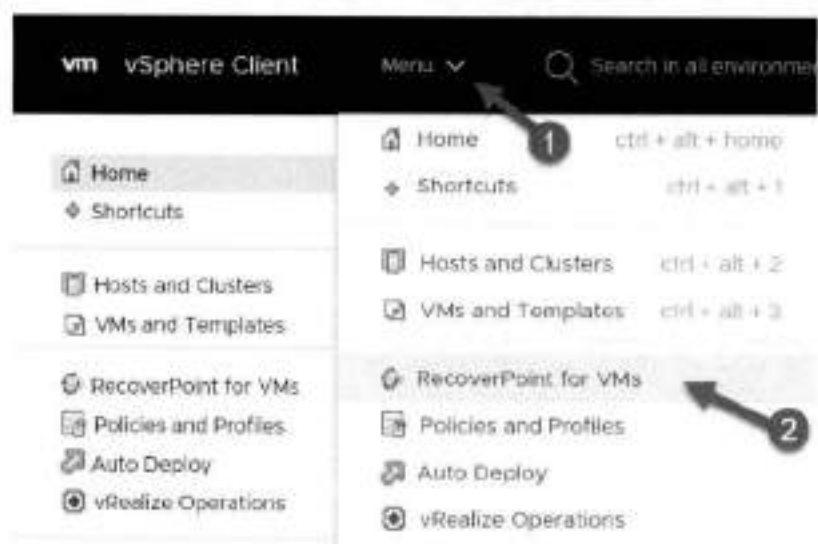
- enable pre-emptive support services.
- display recovery activity reports.
- display performance statistics for consistency groups and vRPA clusters.
- protect VMs with more than one local copy and two remote copies.
- select an existing copy VM when adding a VM to an existing consistency group.

In RecoverPoint for VMs versions prior to 5.3 SP2, when running a vSphere version that supports FLEX, also use the **vSphere FLEX plugin** to monitor system health, components, limits, events, usage, and capacity.

NOTE: On vSphere versions that do not support FLEX, you can only perform these tasks through the *RecoverPoint for VMs RESTful API* at <https://developer.dell.com/apis>.

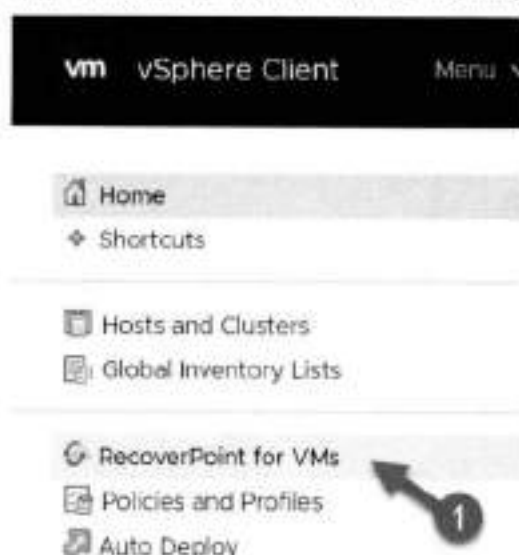
In all other cases, click **LAUNCH VSPHERE CLIENT (HTML5)** to display the RecoverPoint for VMs **vSphere HTML5 plugin**.

3. Display the RecoverPoint for VMs plugin in the vSphere Client:
- Click the **RecoverPoint for VMs** menu item in your main vSphere Client **Menu**.



-OR-

- Click the **RecoverPoint for VMs** menu item in your vSphere Client **Navigator**.



Results

The RecoverPoint for VMs Dashboard is displayed.

Add license

Add a RecoverPoint for VMs license for each vCenter Server that is hosting a protected VM. Adding a license automatically registers the product and enables support.

Prerequisites

Create your license files. To learn more about the types of licenses that are available, see RecoverPoint for VMs licensing.

Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, click **System > Licenses > Add**.

Add a RecoverPoint for VMs License



Access your entitlements in the 'Software Licenses' section of support.dell.com or by clicking the link in the LAC email sent to the email address provided during order entry. Activate your entitlements and download your license (*.lic) files from support.dell.com. Add a license for each vCenter Server that is protecting a VM or hosting a vRPA cluster. To learn more about the types of licenses that are available, see the RecoverPoint for VMs Administrator's Guide.

ⓘ When you add a socket-based license to a system with VM-based licenses, the system converts VM licenses to socket licenses at a ratio of 15 VMs per socket. When the ratio is not an even conversion, the value is rounded up.

Select a license (*.lic) file

CANCEL

ADD LICENSE

The **Add a RecoverPoint for VMs License** screen is displayed.

2. Click the browse icon, select a license file (*.lic), and click **Add License**.

Results

Your license is added to the RecoverPoint for VMs system, and its usage statistics and expiration are displayed in the **RecoverPoint for VMs Licenses** table. When both VM-based and socket-based licenses are added, the VM licenses are converted to socket licenses at a rate of 15 VMs to one socket.

To register your system, an email with your license details is automatically sent to emailalerts@dell.com.

Register for Customer Support

If the automatic registration email was not sent during license addition, or a change was made to your system follow this procedure. By ensuring that Customer Support has up-to-date information on the configuration of your system, you enable Dell to provide you with the most effective support possible.

About this task

Refer to the *RecoverPoint for VMs CLI Guide* for more detailed information.

Steps

1. Create an SSH connection to a vRPA cluster management IP address, and use your RecoverPoint for VMs admin username and password to log into the Boxmgmt CLI and, from there, open the Sysmgmt CLI.
2. In the Sysmgmt CLI, run the **set_registration_params** command, and provide all of the requested information.

Results

You should receive notification that your registration parameters have been successfully configured. Run the **get_registration_params** command if you want to confirm that the information in your system registration is correct.

Protecting VMs

In RecoverPoint for VMs, consistency groups are used to protect virtual machines and replicate virtual machine application data to a consistent point in time. A consistency group is a logical entity that constitutes a container for virtual machines and all of their copies.

Consistency groups can protect many VMs. If this is the first time you are using RecoverPoint for VMs, protect your VMs by creating new consistency groups for them. If you already have RecoverPoint for VMs consistency groups, you can protect your VMs by creating new consistency groups for them, or by adding them to an existing consistency group. You can also create a new copy to protect your production VMs, alongside your existing copy.

For additional information about VM management, including how to unprotect VMs, see *Managing virtual machines*.

NOTE: Protecting a virtual machine with fault tolerance enabled is not supported.

Before protecting your VMs, ensure you have:

- Completed the tasks described in *Before you begin*.
- Powered on the virtual machines that you want to protect.
- Registered all linked vCenter Servers hosting production VMs and copy VMs as described in *Managing the plugin server*.

Topics:

- Protect a VM in a new group
- Protect a VM in an existing group
- Protect multiple VMs in a new group
- Protect multiple VMs in an existing group
- Add a copy to an existing group

Protect a VM in a new group

Protect a virtual machine in a new consistency group.

Steps

1. Connect to the vSphere Client of your production site.
2. Select **VMs and Templates** view.
3. Right-click a powered-on VM, and select **RecoverPoint for VMs > Protect VM...**
The **Protect VM** dialog is displayed.

Protect VM



other_os [Edit Settings](#)

Protected by: Consistency Group: cg_other_os

Production: vRPA Cluster: Site1, Journal Datastore: DEV_RPVE34_Site

Copies: [+ ADD A COPY](#)

vRPA Cluster: Site2 (Remote Cop)	Sync: RPO: 25	Async: seconds	vCenter Server: VM-RP-LAB-VC-28	Target ESXi Cluster: Site 2	Copy Datastore: DEV_RPVE34_Site
----------------------------------	---------------	----------------	---------------------------------	-----------------------------	---------------------------------

[CANCEL](#) [PROTECT](#)

NOTE: In the **Protect VM** dialog, all of the fields are pre-populated with sensible values, so you can safely click **PROTECT** now, and manage the protection policies later, if necessary.

- (Optional) Click **Edit Settings** to change the default VM protection policy.

Hardware

Disk Provisioning: Same as source

Replicate Hardware Changes

VMDKs (4 / 4)

<input checked="" type="checkbox"/> Hard disk 1 - SCSI (0:0)	3 GB
<input checked="" type="checkbox"/> Hard disk 2 - SCSI (0:1)	1 GB
<input checked="" type="checkbox"/> Hard disk 3 - SCSI (0:2)	2 GB
<input checked="" type="checkbox"/> Hard disk 4 - SCSI (0:3)	3 GB

- Disk Provisioning:** Default is **Same as source**. Defines how the copy VMDKs are provisioned; **Same as source**, **Thick provision lazy zeroed**, **Thick provision eager zeroed** or **Thin provision**.
- Replicate Hardware Changes:** Default is **Enabled**. Automatically replicates the hardware settings of all production virtual machines to their copy VMs whenever an image is accessed on the copy VMs. When enabled, RecoverPoint for VMs replicates the virtual machine version, CPU, memory, resource reservations, and network adapter status and type, and MAC addresses (only to remote copy VMs).

NOTE: Replication of SR-IOV Passthrough Adapter is not supported. If the ESXi at a copy does not support the production VM version, no hardware changes are replicated.
- VMDKs:** Displays the number of VMDKs that will be replicated, and their total size. Clear a VMDK check box to exclude the VMDK from replication.

- (Optional) Change the default consistency group protection policies.

Protected by: Consistency Group: cg_other_os

Production: vRPA Cluster: Site1, Journal Datastore: DEV_RPVE34_Site

- Consistency Group:** Default is **cg_<vmname>**. Defines the consistency group name.
- vRPA Cluster:** Defines the vRPA cluster used to replicate and manage the production data to the copies.
- Journal Datastore:** Defines the datastore that RecoverPoint for VMs will automatically provision a **3GB** VMDK on for the production journal. By default, RecoverPoint automatically registers up to 15 datastores for the production and copy journals and automatically selects the datastore with the most free space.

- NOTE:** RecoverPoint for VMs will attempt to create the production journal on the selected datastore. If it cannot, the system will attempt to create the production journal on another registered datastore. If you have more than 15 datastores and would like to register an additional datastore that is not in the list, register the other datastores according to Managing journal datastore registration.

6. (Optional) Change the default copy protection policy.

Update the copy policies:

vRPA Cluster Site2 (Remote Cop	Sync <input checked="" type="checkbox"/> Async RPO 25 seconds	vCenter Server VM-RR-LAB-VC-28	Target ESXi Cluster Site 2	Copy Datastore DEV_RPvE34_Site
--	--	--	--------------------------------------	--

- **vRPA Cluster:** Defines the vRPA cluster used to replicate and manage the production VM data to the storage at this copy.
- **Sync/Async:** Default is **Asynchronous** with **RPO (Recovery Point Objective) of 25 seconds**. The RPO is the point in time to which you are required to recover data, for a specific application, as defined by the organization. RPO defines the maximum lag that is allowed on a link.
- **vCenter Server and ESX Cluster:** Defines the vCenter Server and ESX cluster hosting the copy VMs.
- **Copy Datastore:** Defines the datastore to use for the copy VM data.

7. (Optional) Click the copy's **Advanced Configuration** icon to update the advanced copy policies:



Protect VM



vRPA Cluster Darwin (Remote Cr	Sync <input checked="" type="checkbox"/> Async RPO 25 seconds	vCenter Server VM-RR-Lab-H-134	Target ESXi Cluster Darwin	Copy Datastore DD_StorageUnit_C
--	--	--	--------------------------------------	---

Advanced Copy Configuration

Journal Datastore DD_StorageUnit_C	Journal Size 10 GB
--	------------------------------

Copy VM Creation

Production VM	Copy VM	RESET ALL TO DEFAULT
Deploy_RPCenter_mgtfanc_913	<input checked="" type="radio"/> Automatically create copy VM <input type="radio"/> Manually select copy VM Select a Resource...	

- **Journal Datastore:** Defines the datastore that RecoverPoint for VMs will automatically provision a 10GB VMDK on for the copy journal, unless **Manually select copy VM** is selected as the **Copy VM Creation** method. By default, RecoverPoint automatically registers up to 15 datastores for the production and copy journals and automatically selects the datastore with the most free space.

- NOTE:** RecoverPoint for VMs will attempt to create the copy journal on the selected datastore. If it cannot, the system will attempt to create the copy journal on another registered datastore. If you have more than 15 datastores and would like to register an additional datastore that is not in the list, register the other datastores according to Managing journal datastore registration.

- **Journal Size:** Default is **10GB**. The larger the copy journal, the more history can be saved.
- **Copy VM Creation:** Default is **Automatically create copy VM**. You can:
 - Click **Select a Resource...** and select the ESXi host to host the copy VM.

- Select **Manually select copy VM** > click **Select a VM...** and select a VM from the list.
8. (Optional) For added protection, click **ADD A COPY** to protect the VM with an additional copy:
- Up to two copies can be created during VM protection. For additional protection, Add a copy to an existing group to create more copies.
 - After adding a copy, you can click **Delete Copy** to delete a copy. The last copy cannot be deleted.



9. Click **PROTECT**.

Results

The specified virtual machine is protected, and the group production data starts being replicated to the copy VMs according to the specified policies. If an unregistered ESX cluster, or the VMware Resource Pool of an unregistered ESX cluster was selected to host a copy VM, the unregistered ESX cluster is automatically registered with the specified vRPA cluster, a splitter is installed on every ESXi host in the cluster, and replication is temporarily paused for all relevant VMs while the splitter is being installed.

Next steps


- See Managing VM protection policies for additional VM protection policies.
- See Managing group protection policies for additional group and copy protection policies.
- For additional protection, Add a copy to an existing group to create more copies.

Protect a VM in an existing group

Protect a virtual machine in an existing consistency group.

About this task

The VM that you select to protect will be added to an existing consistency group that is already protecting a VM. For best performance, you should only protect one VM per consistency group.

 **CAUTION: If the image of the VM that you want to protect is larger than the journal size of the copy, the system automatically enters *one-phase distribution mode* upon protection.**

Steps

1. Connect to the vSphere Client of your production site.
2. Select **VMs and Templates** view.
3. Right-click a powered-on VM, and select **RecoverPoint for VMs > Protect VM in Existing Group...**

The **Protect this VM in an Existing Group** dialog is displayed.

Protect this VM in an Existing Group



Test-VM

Edit Settings

Consistency Groups HELP

cg_TVM-1 is selected.

Consistency Group	Production vCenter Server	Production vSphere Cluster	State
cg_TVM-1	VM-RP-Lab-H213 hrz.lab.dell.com	VSPA_Test	Enabled
Copies			
vCenter Server vm-rp-lab-h-213...	Target ESXi Cluster London	Copy Datastore RPVENV2_LOND	
cg_TVM2	VM-RP-Lab-H213 hrz.lab.dell.com	VSPA_Test	Enabled

CANCEL

PROTECT

NOTE: All of the fields are pre-populated with sensible values, so you can safely select the consistency group, and click **PROTECT** now. You can manage the VM protection settings later, if necessary, as described in Managing VM protection policies.

4. Select the consistency group to protect your production VMs.

When a consistency group is selected, the group copies are displayed and you can change the **Target ESXi Cluster** and **Copy Datastore** of the group copy VMs.

5. (Optional) To change the VM creation settings, perform the following actions:

- a. Click . The **Copy VM creation** page appears.
- b. **Copy VM** is set to **Automatically create copy VM** by default. To change the **Copy VM** setting, select the appropriate option. When **Automatically create copy VM** is selected, from the drop-down list, you can select the ESXi host to host the copy VM. When **Manually select copy VM** is selected, from the drop-down list you can select a VM.
- c. To reset all of the **Copy VM Creation** setting values to the default values, click **RESET ALL TO DEFAULT**.
- d. To save your changes, click **BACK**.

6. (Optional) Click **Edit Settings** to change the default VM protection policies.

Hardware

Disk Provisioning Same as source

Replicate Hardware Changes

VMDKs (4/4)

- | | |
|--|------|
| <input checked="" type="checkbox"/> Hard disk 1 - SCSI (0:0) | 5 GB |
| <input checked="" type="checkbox"/> Hard disk 2 - SCSI (0:1) | 1 GB |
| <input checked="" type="checkbox"/> Hard disk 3 - SCSI (0:2) | 2 GB |
| <input checked="" type="checkbox"/> Hard disk 4 - SCSI (0:3) | 3 GB |

See Managing VM protection policies for a detailed description of these VM protection policies, and others that cannot be defined during VM protection.

7. Click **PROTECT**.

Results

The specified virtual machine is protected in the specified consistency group.

- A volume sweep occurs on the newly added VM and a short initialization occurs on all other VMs in the consistency group.
- If an unregistered ESX cluster, or the VMware Resource Pool of an unregistered ESX cluster, was selected to host a copy VM, the unregistered ESX cluster is automatically registered with the specified vRPA cluster, a splitter is installed on every ESX host in the cluster, and replication is temporarily paused for all relevant VMs while the splitter is being installed.
- RecoverPoint for VMs will attempt to create the journals on the selected datastores. If for any reason journal creation fails, the system will attempt to create the journal on another registered datastore.
- If the image of the VMs that you want to protect is larger than the journal size of the copy, the system automatically enters *one-phase distribution mode* upon protection.

Protect multiple VMs in a new group

Protect multiple virtual machines hosted on an ESX cluster in a new consistency group.

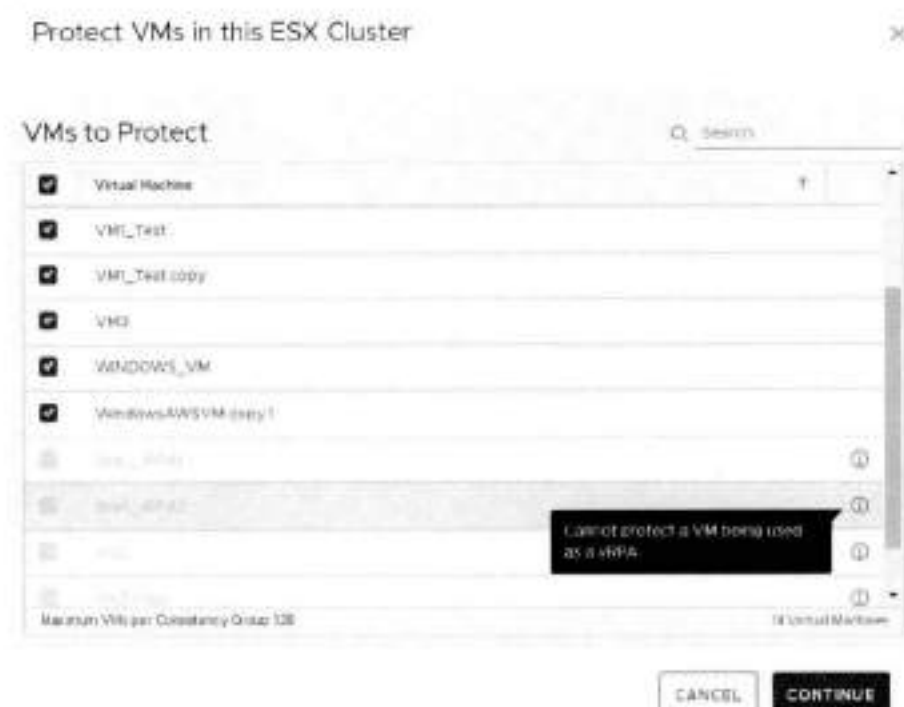
About this task

All of the VMs that you select in the following procedure will be added to a single consistency group. For best performance, you should only protect one VM per consistency group.

Steps

1. Connect to the vSphere Client of your production site.
2. Select **Hosts and Clusters** view.
3. Right-click on an ESX cluster, and select **RecoverPoint for VMs > Protect VMs....**

The **Protect VMs in this ESX Cluster** dialog is displayed.



RecoverPoint for VMs automatically detects if a VM is not able to be protected. Scroll to the bottom of the VM list and click the info icon next to an excluded VM to display the reason for its exclusion.

4. Select the VMs that you want to protect, and click **CONTINUE**.

The **Protect VMs** dialog is displayed.

Protect VMs

VM2	Edit Settings	KATE_WINDOWS	Edit Settings	other_os	Edit Settings
WINDOWS_VM	Edit Settings	VM1_Test.copy	Edit Settings	Deploy_RPCenter_mur	Edit Settings
VM1_Test	Edit Settings	LINUX_VM	Edit Settings	WindowsAWSVM.copy1	Edit Settings

Protected by

Production

Consistency Group
cg_9-vm5

vRPA Cluster
Site1

Journal Datastore
DEV_RPVE34_Site

Copies

+ ADD A COPY

vRPA Cluster
Site1 (Local Copy)

Sync
RPO 25 seconds

Async

vCenter Server
VM-RR-LAB-VC-25

Target ESXi Cluster
Site1

Copy Datastore
DEV_RPVE34_Site

Space will soon be insufficient
271 GB space is required

CANCEL

PROTECT

- NOTE:** In the **Protect VM** dialog, all of the fields are pre-populated with sensible values, so you can safely click **PROTECT** now, and manage the protection settings later, if necessary, as described in [Managing VM protection policies](#) and [Managing group protection policies](#).

5. (Optional) Click **Edit Settings** to change the default VM protection policies.

Hardware

Disk Provisioning Same as source

Replicate Hardware Changes

VMDKs (4 / 4)

<input checked="" type="checkbox"/> Hard disk 1 - SCSI (0:0)	5 GB
<input checked="" type="checkbox"/> Hard disk 2 - SCSI (0:1)	1 GB
<input checked="" type="checkbox"/> Hard disk 3 - SCSI (0:2)	2 GB
<input checked="" type="checkbox"/> Hard disk 4 - SCSI (0:3)	3 GB

See [Managing VM protection policies](#) for a detailed description of these VM protection policies, and others that cannot be defined during VM protection.

6. (Optional) Change the default consistency group and production protection policies.

Protected by

Production



Consistency Group
cg_other_os

vRPA Cluster
Site1

Journal Datastore
DEV_RPVE34_Site

See [Managing group protection policies](#) for a detailed description of these group protection policies, and others that cannot be defined during VM protection.

7. (Optional) Change the default copy protection policies.

vRPA Cluster Site2 (Remote Cop -	Sync <input checked="" type="checkbox"/> Async RPO 25 seconds	vCenter Server VM-RP-LAB-VC-28 -	Target ESXi Cluster Site 2 -	Copy Datastore DEV_RPVE34_Site -	 
-------------------------------------	--	-------------------------------------	---------------------------------	-------------------------------------	---

See Managing group protection policies for a detailed description of these copy protection policies, and others that cannot be defined during VM protection.

8. (Optional) Update the advanced copy policies, see Protect a VM in a new group for details.
9. (Optional) For added protection, click **ADD A COPY** to protect the VMs with an additional copy:
 - Up to two copies can be created during VM protection. For additional protection, use Add a copy to an existing group to create additional copies.
 - After adding an additional copy, you can click the **Delete Copy** icon to delete a copy. The last copy cannot be deleted.



10. Click **PROTECT**.

Results

The specified virtual machines are protected in a new consistency group.


- If an unregistered ESX cluster, or the VMware Resource Pool of an unregistered ESX cluster, was selected to host a copy VM, the unregistered ESX cluster is automatically registered with the specified vRPA cluster, a splitter is installed on every ESXi host in the cluster, and replication is temporarily paused for all relevant VMs while the splitter is being installed.
- RecoverPoint for VMs will attempt to create the journals on the selected datastores. If for any reason journal creation fails, the system will attempt to create the journal on another registered datastore.

Protect multiple VMs in an existing group

Protect multiple virtual machines hosted on an ESX cluster in an existing consistency group.

About this task

All of the VMs that you select in the following procedure will be added to a single consistency group. For best performance, you should only protect one VM per consistency group.

 **CAUTION:** If the image of the VMs that you want to protect is larger than the journal size of the copy, the system automatically enters *one-phase distribution mode* upon protection.

Steps

1. Connect to the vSphere Client of your production site.
2. Select **Hosts and Clusters** view.
3. Right-click on an ESX cluster, and select **RecoverPoint for VMs > Protect VMs in Existing Group...**

The **Protect VMs in an Existing Group** dialog is displayed.

Protect VMs in an Existing Group



RecoverPoint for VMs automatically detects if a VM is not able to be protected. Scroll to the bottom of the VM list and click the info icon next to an excluded VM to display the reason for its exclusion.

4. Select the VMs that you want to protect, and click **CONTINUE**.

The **Protect these VMs in an Existing Group** dialog is displayed.


Protect this VM in an Existing Group

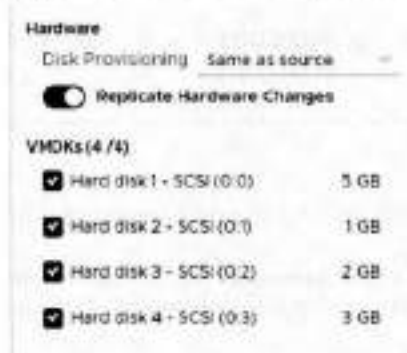


NOTE: All of the fields are pre-populated with sensible values, so you can safely select the consistency group, and click **PROTECT** now. You can manage the VM protection settings later, if necessary, as described in *Managing VM protection policies*.

5. Select the consistency group to protect your production VMs.

When a consistency group is selected, the group copies are displayed and you can change the **Target ESXi Cluster** and **Copy Datastore** of the group copy VMs.

6. (Optional) To change the VM creation settings, perform the following actions:
 - a. Click . The **Copy VM creation** page appears.
 - b. **Copy VM** is set to **Automatically create copy VM** by default. To change the **Copy VM** setting, select the appropriate option. When **Automatically create copy VM** is selected, from the drop-down list, you can select the ESXi host to host the copy VM. When **Manually select copy VM** is selected, from the drop-down list you can select a VM.
 - c. To reset all of the **Copy VM Creation** setting values to the default values, click **RESET ALL TO DEFAULT**.
 - d. To save your changes, click **BACK**.
7. (Optional) Click **Edit Settings** to change the default VM protection policies.



See *Managing VM protection policies* for a detailed description of these VM protection policies, and others that cannot be defined during VM protection.

8. Click **PROTECT**.

Results

The specified virtual machines are protected in the specified consistency group.

- A volume sweep occurs on the newly added VMs and a short initialization occurs on all other VMs in the consistency group.
- If an unregistered ESX cluster, or the VMware Resource Pool of an unregistered ESX cluster, was selected to host a copy VM, the unregistered ESX cluster is automatically registered with the specified vRPA cluster, a splitter is installed on every ESXi host in the cluster, and replication is temporarily paused for all relevant VMs while the splitter is being installed.
- RecoverPoint for VMs will attempt to create the journals on the selected datastores. If for any reason journal creation fails, the system will attempt to create the journal on another registered datastore.
- If the image of the VMs that you want to protect is larger than the journal size of the copy, the system automatically enters *one-phase distribution mode* upon protection.

Add a copy to an existing group

For added protection, add another copy to an existing consistency group.

Steps

1. In the RecoverPoint for VMs plugin for vSphere Client, select **Protection > Consistency Groups**.
2. Select a group, and click the more group actions button (...) > **Add a copy**.
3. In the **Add a Copy** dialog:

Add a Copy to 'MyGroup'

X

vRPA Cluster WELL (Remote Co.)	Sync <input type="radio"/> Async <input checked="" type="radio"/> RPO 25 seconds	vCenter Server VM-RP-Lab-VC-15r	Target ESXi Cluster RPVENv8_Site3	Copy Datastore vsanDatastore (1.2)
-----------------------------------	---	------------------------------------	--------------------------------------	---------------------------------------

NOTE: All of the fields are pre-populated with sensible values, so you can safely click **ADD COPY** now, and manage the protection policies later, if necessary.

- (Optional) Change the default copy protection policy.

Update the copy policies:

- vRPA Cluster:** Defines the vRPA cluster used to replicate and manage the production VM data to the storage at this copy.
- Sync/Async:** Default is **Asynchronous** with **RPO (Recovery Point Objective) of 25 Seconds**. The RPO is the point in time to which you are required to recover data, for a specific application, as defined by the organization. RPO defines the maximum lag that is allowed on a link.
- vCenter Server and ESX Cluster:** Defines the vCenter Server and ESX cluster hosting the copy VMs.
- Copy Datastore:** Defines the datastore to use for the copy VM data.

- (Optional) Update the advanced copy policies, see *Protect a VM in a new group* for details.
- Click **ADD COPY**.

Results

A copy is added to the consistency group, and the group production data starts being replicated to the copy VMs according to the specified policies. If an unregistered ESX cluster, or the VMware Resource Pool of an unregistered ESX cluster, was selected to host a copy VM, the unregistered ESX cluster is automatically registered with the specified vRPA cluster, a splitter is installed on every ESXi host in the cluster, and replication is temporarily paused for all relevant VMs while the splitter is being installed.

Next steps

See *Managing group protection policies* for additional copy protection policies.

Monitoring protection

After protecting your VMs, use the **Dashboard** to monitor the number and status of protected VMs, consistency groups, group sets, vRPA clusters, system alerts and limits, as well as your RecoverPoint for VMs license usage.

Topics:

- The RecoverPoint for VMs Dashboard
- Monitor system alerts
- Monitor system events
- Monitor system limits
- Monitor system components
- Monitor group and copy protection

The RecoverPoint for VMs Dashboard

The **RecoverPoint for VMs Dashboard** presents a high-level overview of the RecoverPoint for VMs system to help you analyze and monitor your system health.



Figure 1. The RecoverPoint for VMs Dashboard

The **Dashboard** is displayed each time you log into RecoverPoint for VMs. Use the **Dashboard** to monitor the status of your system licenses, limits, alerts, protected VMs, consistency groups, group sets, and recovery activities.

In every **Dashboard widget**, you can:

- **Click the help icon** to display more detailed information on the widget system component or activity.



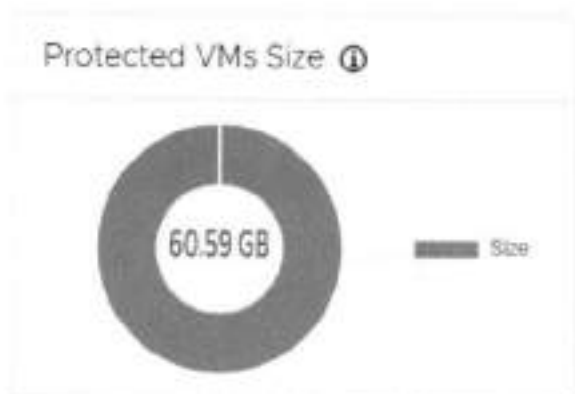
- **Click a status in a legend** (the colored bar or the label) to filter what is shown in the chart and display the number of system components with the status that you clicked. To clear the filter, click the status again.
- **Click a color of a status in a chart** to go to the relevant system component or activity screen, and display only the system component or activity in the clicked status.

The **Dashboard** in RecoverPoint for VMs 5.3 SP2 and later versions contains the following **widgets** that do not exist in previous RecoverPoint versions.

- **Limits**




- Displays the number and status of the limits imposed on system components like consistency groups, splitters, and vRPA clusters.
 - A system component's limit status can be 'Critical', 'Warning', or 'OK'.
 - Click a status in the chart to display all system components with that status in the **Monitoring > System Limits** screen.
 - For more information, see how to Monitor system limits.
- **Protected VMs Size**



- Displays the total size (in GB) of all protected VMs on this vCenter Server.
 - For more information, see Protecting VMs.
- **Unprotected VMs**


Unprotected VMs ⓘ



- Displays the percentage of unprotected VMs.
- When you hover over the widget, the exact number of unprotected VMs appears. This is a warning for the user that this number of VMs are unprotected and have to be protected to make data protection available.
- The amount of unprotected VMs = total VMs - protected VMs - vRPAs - RPC - replica/shadow VMs.

• vRPA Clusters


vRPA Clusters ⓘ



- Displays the number and status of all registered vRPA clusters on all registered and linked vCenter Servers.
- The status of a vRPA cluster can be 'Error', 'Warning', or 'OK'.
- Click a status in the chart to display all vRPA clusters with that status in the **System > Administration > vRPA Clusters** screen.
- For more information, see *Managing vRPA clusters* and *Collecting logs from vRPA clusters*.

• License Usage

License Usage ⓘ



- Displays the number of protected sockets out of the total number of licensed sockets.
- A system license's usage status can be 'Trial', 'OK', or 'Violated'.
- Click a status in the chart to display all licenses with that status in the **System > Licenses** screen.
- For more information, see *Managing your RecoverPoint for VMs licenses*.

The **Dashboard** in RecoverPoint for VMs versions prior to 5.3 SP2 also contains these **widgets**:

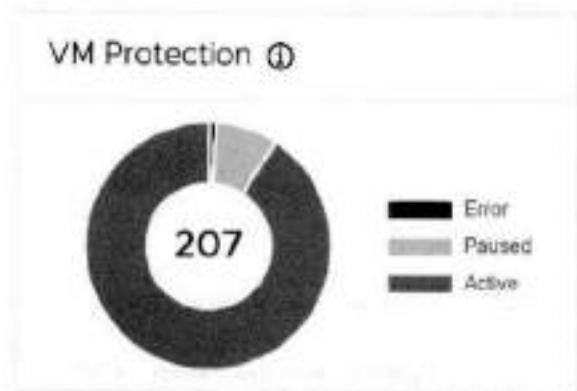
NOTE: RecoverPoint for VMs versions prior to 5.3 SP2 do not contain the **Limits**, **vRPA Clusters**, **License Usage**, and **Protected VMs Size** widgets.

- **Alerts**



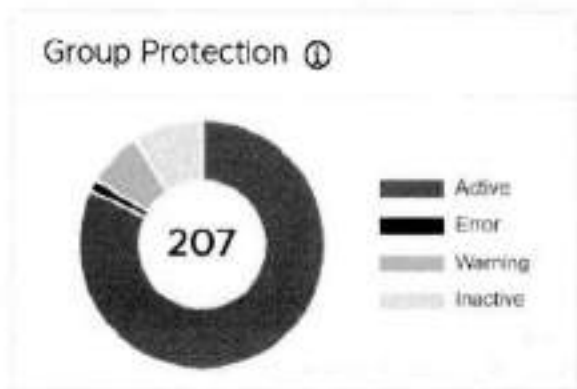
- Displays the number and types of alerts in the system.
- An alert's type can be 'Error' or 'Warning'.
- Click an alert type in the chart to display all alerts of that type in the **Monitoring > Alerts** screen.
- For more information, see Monitor system alerts.

- **VM Protection**



- Displays the number and status of protected VMs on the vCenter Server that you are connected to, or a registered vCenter Server linked to the vCenter that you're connected to.
- A protected VM's status can be 'Active', 'Error' or 'Paused'.
- Click a status in the chart to display all protected VMs with that status in the **Protection > Protected VMs** screen.
- For more information, see Protecting VMs.

- **Group Protection**



- Displays the number and status of all consistency groups in the system.

- o A consistency group's status can be **Active**, **Inactive**, **Warning** or **Error**.
- o Click a status in the chart to display all groups with that status in the **Protection > Consistency Groups** screen.
- o For more information, see *Protecting VMs*.

- **Group Recovery Activities**



- o Displays the number and status of all consistency group recovery activities.
- o The status of an activity can be **Error**, **Action Needed**, or **In Progress**.
- o Click a status in the chart to display all groups with that status in the **Recovery Activities > Consistency Groups** screen.
- o For more information, see *Monitoring recovery activities*.

- **Group Set Recovery Activities**



- o Displays the number and status of all group set recovery activities.
- o The status of an activity can be **Error**, **Action Needed**, or **In Progress**.
- o Click a status in the chart to display all group sets with that status in the **Recovery Activities > Group Sets** screen.
- o For more information, see *Monitoring recovery activities*.

Monitor system alerts

System alerts is a mechanism that enables vRPA clusters to send events about system components in real time. Monitor system alerts to troubleshoot your RecoverPoint for VMs environment.

Prerequisites

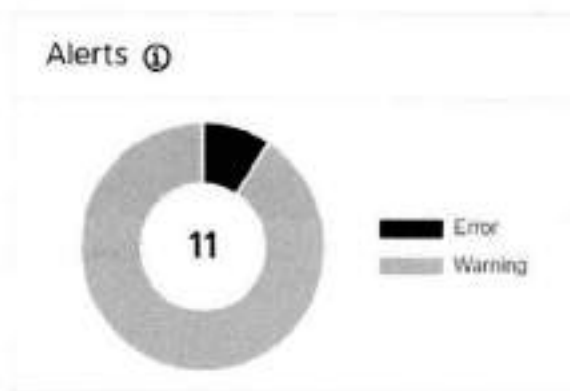
Ensure you have performed the procedure for *Configuring email alerts and reports* and have added valid licenses as described in *Managing your RecoverPoint for VMs licenses*.

About this task

To monitor your system alerts, click **Monitoring > Alerts**. A system alert's type can be **Warning** or **Error**.

Name	Category	Severity	Description
Veeva	System	Error	You are using the default password for the administrator. To change the default password, log in to the UI.
	System	Warning	You are using an unlicensed trial version for non-commercial use only of Dell EMC RecoverPoint for VMs.
Delphide	Consistency g...	Warning	The virtual RPA is running on the same ESX as the VM it is replicating. It is recommended not to replicate th...
Beltaine	RPA	Warning	At least two virtual RPAs are running on same ESX. It is recommended to run only one virtual RPA on a...
Vitani	Consistency g...	Warning	At least one virtual RPA is running on the same ESX as the VM it is replicating. It is recommended not t...
Alerts per page	25		3 Alerts

NOTE: You can also monitor your system alerts in the The RecoverPoint for VMs Dashboard.



Monitor system events

Monitor system events to troubleshoot your RecoverPoint for VMs environment.

An event is a notification that a change has occurred in the state of a system component. In some cases, the change indicates an error or warning condition for a system component. Multiple events can occur simultaneously on a single component and a single incident can generate multiple events across multiple system components.

By default, the following information is displayed for every event in the **Event Logs**:

- **Level**, which can be: **Info**, **Warning**, or **Error**.
- **Scope**, which can be: **Normal**, **Detailed**, or **Advanced**.
- **Time** and date that the event log was generated.
- **vRPA Cluster** reporting the event.
- **Event ID** that allows the event to be excluded from the events log using the event logs filter.
- **Topic**, which can be: **Splitter**, **Consistency Group**, **Management**, **Cluster**, **RPA**, or **Array**.
- **Summary** of the event.

To monitor your system events, click **Monitoring > Event Logs**.

Event Logs

Search 4 3

Warnings and errors only from the last 24 hours

Last Updated: Aug. 2021 3:00:00 PM

Level	Scope	Time	vRPA Cluster	Event ID	Topic	Summary
Transient Error	Root Cause	Jul 7, 2021 4:53:54 PM	T3-Cluster	9038	RPA	Brief group(s) journal volumes accessibility error es...
Transient Error	Root Cause	Jul 7, 2021 4:54:11 PM	NY-Cluster	9034	RPA	The data path between all RPAs and all spinters (and storage) was briefly unavailable but the problem has been corrected.
Transient Error	Root Cause	Jul 7, 2021 4:54:11 PM	NY-Cluster	9030	RPA	Brief group(s) journal volumes accessibility error es...
Error	Normal	Jul 7, 2021 5:03:04 PM	NY-Cluster	4003	Consistency Group	Group capabilities problem
Error	Normal	Jul 7, 2021 5:21:04 PM	NY-Cluster	4003	Consistency Group	Group capabilities problem
Error	Normal	Jul 7, 2021 5:59:31 PM	NY-Cluster	4003	Consistency Group	Group capabilities problem
Error	Normal	Jul 7, 2021 5:59:45 PM	NY-Cluster	4005	Consistency Group	Group capabilities problem
Error	Normal	Jul 7, 2021 6:00:04 PM	NY-Cluster	4002	Consistency Group	Group capabilities problem
Error	Normal	Jul 7, 2021 6:00:45 PM	NY-Cluster	4002	Consistency Group	Group capabilities problem
Error	Normal	Jul 7, 2021 6:29:06 PM	NY-Cluster	4005	Consistency Group	Group capabilities problem

Items per page: 20 1 10 Event Logs

Figure 2. RecoverPoint for VMs event logs

1. Note the total number of events in the event logs.
2. Use the table controls to move to the next page, a previous page, or control the number of events displayed per page.
3. Click the **Event Filter** to control which events are displayed in the **Event Logs** and which are hidden.

Filter APPLY CANCEL

- > vRPA Cluster
- > Time Range
- > Topics
- > Scope
- > Level
- > Event IDs to exclude

NOTE: Click **APPLY** after changing the event filter settings.

- Click **vRPA Cluster** to select the events for a specific vRPA cluster to display. By default, the events of all vRPA clusters are displayed.
- Click **Time Range** to select the events of a specific time period to display. Select **Unbound** to display all events. Display events based on your local time (the default) or GMT.

Filter

APPLY

CANCEL

> VRPA Cluster

Select Cluster: ALL ▾

> Time Range

From

Unbound

07/17/2021  06:22 PM 

To

Unbound

07/18/2021  06:22 PM 

Local time GMT

- Click **Topics** to hide or display events for specific system components. The event topic can be: **Splitter**, **Consistency Group**, **Management**, **Cluster**, **RPA**, or **Array**.
- Click **Scope** to hide or display logs of specific event scope. The event scope can be: **Normal**, **Detailed**, or **Advanced**.
- Click **Level** to hide or display events of a specific level. The event level can be **Info**, **Warning**, or **Error**.
- Click **Event IDs to exclude** to select the events to exclude from display in the events log.

Filter

APPLY

CANCEL

> vRPA Cluster

Select Cluster ALL

> Time Range

From

Unbound

04/07/2022 12:34 PM

To

Unbound

04/08/2022 12:34 PM

Local time GMT

> Topics

Splitter

Consistency Group

Management

Cluster

RPA

Array

> Scope

Root Cause Normal Advanced

> Level

Error Warning Info

Info level is available only if a single cluster is selected.

> Event IDs to exclude

Add an event ID

- Note the date and time that the **Event Logs** were **Last Updated** and use the **Refresh** icon to update the **Event Logs**.
- Hover over the **Summary** of an event with an ellipsis (...) after it, to display hidden text.

While troubleshooting:

- Use the search bar to display only events that include specific text.

Search



- Click the **Clear Filters** button to easily clear all event filters.



- Click an arrow to expand an event and display the event **Description** and **Details**.

System Events

Q Search



Warnings and errors only from the last 24 hours

Last updated Apr 6, 2022 6:50:23 PM

	Level	Scope	Time	vRPA Cluster	Event ID	Topic	Summary
▼	Warning	Normal	Apr 6, 2022 ...	local-cluster	5076	Splitter	Splitter has restarted
	Description Splitter has restarted(Splitter(s) [REDACTED])						
>	Warning	Normal	Apr 6, 2022 ...	local-cluster	5076	Splitter	Splitter has restarted
>	Warning	Normal	Apr 6, 2022 ...	local-cluster	4008	Consistency ...	Pausing data transf...
>	Warning	Normal	Apr 6, 2022 ...	local-cluster	4008	Consistency ...	Pausing data transf...
>	Error	Normal	Apr 6, 2022 ...	local-cluster	4009	Consistency ...	Pausing data transf...
>	Warning	Normal	Apr 6, 2022 ...	local-cluster	4001	Consistency ...	Minor problem in gr...
>	Error	Normal	Apr 6, 2022 ...	local-cluster	4009	Consistency ...	Pausing data transf...
>	Warning	Normal	Apr 6, 2022 ...	local-cluster	5076	Splitter	Splitter has restarted
>	Warning	Normal	Apr 6, 2022 ...	local-cluster	5076	Splitter	Splitter has restarted
>	Warning	Normal	Apr 6, 2022 ...	local-cluster	5076	Splitter	Splitter has restarted

Items per page 20

1 / 10

267 System Events

Monitor system limits

Monitor the limits imposed on your RecoverPoint for VMs system and system components to troubleshoot your system.

NOTE: System limits are imposed by your RecoverPoint for VMs licenses, see [Managing your RecoverPoint for VMs licenses](#) for more information.

To monitor the state of your system limits, click **Monitoring > System Limits**. The **System Limits** screen displays the limits imposed on a system, or on consistency groups, vRPA clusters, or splitters in a system.

A system component's limit status can be **Critical**, **Warning**, or **OK**. Ensure an **OK** status is displayed for all of your system limits.

System Limits

Q Search



Consistency Groups **System** Splitters vRPA Clusters

Status	Description	↑	vRPA Cluster	Current Value of Limit
● OK	Number of VMs protected per vC...			2 out of 8192
● OK	Number of vRPA clusters		local cluster	1 out of 5
● OK	Number of vRPA clusters connect...			1 out of 50

Items per page 20

3 System limits

Figure 3. System limits

System Limits

Q Search



Consistency Groups **System** Splitters vRPA Clusters

Status	Description	↑	vRPA Cluster	Consistency Group	Current Value of Limit
● OK	Lag (in seconds)		Seq0	rg_Win5	1 out of 25
● OK	Lag (in seconds)		Seq1	rg_Win1	0 out of 25
● OK	Lag (in seconds)		Seq1	rg_Win1	1 out of 25
● OK	Number of non-production copies		Seq1	rg_Win1	2 out of 4
● OK	Number of non-production copies		Seq1	rg_Win5	1 out of 4
● OK	Number of VMs		Seq1	rg_Win1	1 out of 128
● OK	Number of VMs		Seq1	rg_Win5	1 out of 128

Items per page 20

7 Consistency Groups limits

Figure 4. Consistency group limits

System Limits



Consistency Groups System Splitters **VRPA Clusters**

Status	Description	#	VRPA Cluster	Current Value of Limit
● OK	Number of consistency groups	2	Site1	2 out of 512 groups
● OK	Number of ESX clusters	1	Site2	1 out of 8 clusters
● OK	Number of ESX clusters	1	Site1	1 out of 8 clusters
● OK	Number of ESX hosts with a splitter	2	Site2	2 out of 256
● OK	Number of ESX hosts with a splitter	2	Site1	2 out of 256
● OK	Number of protected VMDKs	4	Site3	4 out of 4096 protected VMDKs
● OK	Number of protected VMDKs	4	Site1	4 out of 4096 protected VMDKs
● OK	Number of protected VMs	2	Site2	2 out of 1024 protected VMs
● OK	Number of protected VMs	2	Site1	2 out of 1024 protected VMs
● OK	Number of vCenter Servers	1	Site1	1 out of 4 VCs
● OK	Number of vCenter Servers	1	Site2	1 out of 4 VCs

Items per page: 20

11 VRPA Clusters Limits

Figure 5. VRPA cluster limits

System Limits



Consistency Groups System **Splitters** VRPA Clusters

Status	Description	#	Splitter	Current Value of Limit
● OK	Number of VRPA clusters attached to splitter	1	Splitter1	1 out of 128
● OK	Number of VRPA clusters attached to splitter	1	Splitter2	1 out of 128
● OK	Number of VRPA clusters attached to splitter	1	Splitter3	1 out of 128
● OK	Number of VRPA clusters attached to splitter	1	Splitter4	1 out of 128
● OK	Total number of VMDKs attached to splitter	4	Splitter1	4 out of 25000
● OK	Total number of VMDKs attached to splitter	4	Splitter2	4 out of 25000
● OK	Total number of VMDKs attached to splitter	6	Splitter3	6 out of 25000
● OK	Total number of VMDKs attached to splitter	6	Splitter4	6 out of 25000

Items per page: 20

8 Splitter Limits

Figure 6. Splitter limits

① **NOTE:** You can also monitor your system **Limits** in the The RecoverPoint for VMs Dashboard.



Monitor system components

Monitor RecoverPoint for VMs system components to better understand and troubleshoot your RecoverPoint for VMs environment.

To monitor the state of your system components, click **Monitoring > Components**. In the **System Components** screen, ensure an OK status is displayed next to each of your RecoverPoint for VMs system components.

System Components

Search ⌵

Status	Name	Type	vRPA Cluster	Version
● OK	vRPA1	vRPA	Site2	vRPA: 5.3 SP2(m.240)
● OK	vRPA2	vRPA	Site2	vRPA: 5.3 SP2(m.240)
● OK	splitter1 - VMware vSphere ESX Host	ESX Splitter	Site2	SPLITTER: 5.3 SP2(m.222) JAM: 5.3 SP2(m.180)
● OK	splitter2 - VMware vSphere ESX Host	ESX Splitter	Site2	SPLITTER: 5.3 SP2(m.222) JAM: 5.3 SP2(m.180)
● OK	vRPA2	vRPA	Site1	vRPA: 5.3 SP2(m.240)
● OK	vRPA1	vRPA	Site1	vRPA: 5.3 SP2(m.240)
● OK	splitter1 - VMware vSphere ESX Host	ESX Splitter	Site1	SPLITTER: 5.3 SP2(m.222) JAM: 5.3 SP2(m.180)
● OK	splitter2 - VMware vSphere ESX Host	ESX Splitter	Site1	SPLITTER: 5.3 SP2(m.222) JAM: 5.3 SP2(m.180)

Items per page: 25 ⌵ 8 System Components

ⓘ **NOTE:** You can also monitor the state of **vRPA clusters** and **Components** in the The RecoverPoint for VMs Dashboard.



Monitor group and copy protection

Monitor the status of replication for consistency groups and copies, when managing or troubleshooting your system.

Steps

1. Select **Protection > Consistency Groups**.
2. Expand a group.
3. Note the **Transfer Status (1)** and **State (2)** of each consistency group and the **Status (3)** of each copy.



Results

- The **Transfer Status (1)** of a consistency group can be:
 - **Active:** Data is being transferred to a copy.
 - **Initializing:** A copy is being initialized: volume sweep, short init, or full sweep.
 - **High Load:** The system enters a temporary high-load state while data is being transferred to a copy, when the journal is full and cannot accept new writes. The system attempts to resolve the high-load state without user action.
 - **Paused by System:** System paused replication so data is not being transferred. If this state occurs for long periods of time, check the system alerts and events in the **Dashboard** for more information.
 - **Error:** An error has occurred.
 - **Permanent High Load:** The system enters a permanent high-load state while data is being transferred to a copy. A permanent high-load can occur after a temporary high-load. The system pauses replication and waits for user action.
 - **Paused:** User paused replication so data is not being transferred to a copy.
 - **Disabled:** User disabled a copy so data is not being transferred.
- The **State (2)** of a consistency group can be:
 - **Enabled:** A group is enabled for replication.
 - **Failed over:** A multi-copy group has completed temporary failover.
 - **Being recovered:** A group is in the process of recovering production.
 - **Partially suspended:** Some of a group's copies have been momentarily suspended while being upgraded.
 - **Suspended:** All of a group's copies have been momentarily suspended while being upgraded.
 - **Disabled:** A group is disabled for replication.
- The **Status (3)** of a copy can be:
 - **OK:** Data can be transferred to the copy.
 - **Initializing:** A copy is being initialized: volume sweep, short init, or full sweep.
 - **High Load:** A copy enters a temporary high-load state while data is being transferred to the copy, when the journal is full and cannot accept new writes. The system attempts to resolve the high-load state for the copy without user action.
 - **Paused by System:** System paused replication so data is not being transferred to a copy. If this state occurs for long periods of time, check the system alerts and events in the **Dashboard** for more information.
 - **Error:** An error has occurred on the copy.
 - **Permanent High Load:** A copy enters a permanent high-load state while data is being transferred to the copy. A permanent high-load can occur after a temporary high-load. The system pauses replication to the copy and waits for user action.
 - **Paused:** User paused replication so data is not being transferred to a copy.
 - **Disabled:** User disabled a copy so data is not being transferred.

VM automation and orchestration

RecoverPoint for VMs provides the following features that automate and orchestrate the recovery of your copy VMs.

Topics:

- Create a bookmark
- Automation
- Orchestration

Create a bookmark

Label a snapshot of a virtual machine, a consistency group, or a group set, for identification during testing and recovery. RecoverPoint for VMs creates crash-consistent snapshots.

About this task

Creating a bookmark on a protected VM creates a bookmark on all copy VMs of the group containing the protected VM. Creating a bookmark on a consistency group that is part of a group set creates a bookmark on all copy VMs of all groups in the group set.

Steps

1. In the RecoverPoint for VMs plugin for vSphere Client:
 - To bookmark a snapshot of a consistency group, click **Protection > Consistency Groups**.
 - To bookmark a snapshot of a protected virtual machine, click **Protection > Protected VMs**.
 - To bookmark a snapshot of a group set, click **Protection > Group Sets**.
2. Select the consistency group, protected VM, or group set that you want to bookmark.
3. Click **BOOKMARK**.
4. In the **Bookmark** dialog:

Bookmark VM 'vm2'

Bookmark
MyBookmark

Snapshot Type
Crash-consistent

Snapshot Consolidation Policy
 Consolidate snapshot

This snapshot must survive daily consolidations

CANCEL CREATE BOOKMARK

- **Bookmark:** Enter a name for the snapshot. The bookmark is the name that is used to identify the snapshot during testing and recovery.

- **Snapshot Type:** Default is **crash-consistent**. Change this value to application-consistent only if you know the snapshot to be application consistent. Selecting **application-consistent** does not create an application-consistent snapshot, it only labels the snapshot as known to be application-consistent.
- **Snapshot Consolidation Policy**
 - **Consolidate snapshot:** Default is **disabled**.
 - **This snapshot must survive daily, weekly, or monthly consolidations:** Default is **daily**.
 - **Daily:** Snapshot survives daily consolidations but is consolidated weekly and monthly.
 - **Weekly:** Snapshot survives daily and weekly consolidations but is consolidated monthly.
 - **Monthly:** Snapshot survives daily, weekly, and monthly consolidations.

5. Click **CREATE BOOKMARK**.

Results

A crash-consistent snapshot is created for the specified VM, group, or group set, with the specified label and consolidation policy.

If the bookmark was created on a:

- Protected VM, the system creates a bookmark on all copy VMs of the group containing the protected VM.
- Consistency group that is part of a group set, the system creates a bookmark on all copy VMs of all groups in the group set.

Next steps

To display bookmarks, go to **Protection > Consistency Groups**, expand a group, and select **Snapshots** from the copy commands.

Automation

This section describes the RecoverPoint for VMs features for automating the replication of virtual machines and VMDKs. VM automation can be defined when protecting VMs, or later through the VM protection policy.

To configure VM automation after protection, select **Protection > Protected VMs**, select a VM, and click **PROTECTION POLICY**.

Protection Policy of VM 'Windows' X

Disk Provisioning Same as source

Replicate VM hardware changes

Replicate M/C addresses to the local copy

Automatically protect newly added VMDKs

Protected VMDKs

<input checked="" type="checkbox"/>	Protected VMDK	Path	Size
<input checked="" type="checkbox"/>	Hard disk 1	SCSI (00)	20 GB
<input checked="" type="checkbox"/>	Hard disk 2	SCSI (01)	200 MB

Automatic protection of newly added VMDKs

Define whether or not VMDKs that are added to a protected VM should automatically be protected.

About this task

By default, all newly added VMDKs are automatically protected. Use this procedure to change the default behavior.

Steps

1. Click **Protection > Protected VMs**.
2. Select the virtual machine for which you want to disable the automatic protection of any newly added VMDKs, in the future.
3. Click **PROTECTION POLICY**.
4. Disable or enable **Automatically protect newly added VMDKs**.
5. Click **UPDATE POLICY**.

Results

The protection policy is updated and RecoverPoint for VMs will use the new policy the next time a VMDK is added to this VM.

Provisioning copy VMDKs

Define the way copy VMs are provisioned, per consistency group.

About this task

By default, copy VMDKs are provisioned **Same as source**. Use this procedure to change the default behavior.

Steps

1. Click **Protection > Protected VMs**.
2. Select a protected VM.
3. Click **PROTECTION POLICY**.
4. In the **Disk Provisioning** drop-down, select **Same as source**, **Thick provision lazy zeroed**, **Thick provision eager**, or **Thin provision**.
5. Click **UPDATE POLICY**.

Results

Newly added copy VMDKs are provisioned according to the specified settings. Copy VMDKs that were already provisioned will not be re-provisioned, and will keep the provisioning method defined during initial protection.

Excluding a VMDK from replication

Include or exclude protected VMDKs from replication.

About this task

Protected VMs containing non-persistent VMDKs cannot be replicated. They should be excluded from replication or their VMDK type should be changed through vSphere Client.

- Excluded VMDKs are not replicated, but their corresponding copy VMDKs are not deleted. Excluded copy VMDKs are created at the copy, but writes going to excluded VMDKs are not replicated to their copy VMDKs.
- Copy VMDKs are created for any production VMDKs, both included and excluded. For most efficient use of storage resources, ensure that disk provisioning is configured as **Thin provision** (or **Same as source**, if production is thin provisioned) before adding the VMDK or upon VM protection.
- Changing the disk type of a non-persistent VMDK to a persistent VMDK does not automatically include the VMDK in replication, even if **Automatically protect newly added VMDKs** is enabled.

Steps

1. Click **Protection > Protected VMs**.
2. Select the VM whose VMDKs you want to exclude from replication.
3. Click **PROTECTION POLICY**.
4. Clear the check box next to each **Protected VMDK** that you want to exclude from replication.
5. Click **UPDATE POLICY**.

Results

In the future, the excluded VMDKs are not displayed in the list of snapshots that you can select when Recovering VMs from a previous point in time, even when recovering snapshots from a time previous to the VMDK removal.

Automatic replication of VM hardware changes

Enables or disables the automatic replication of hardware changes made to a protected VM.

About this task

By default any hardware changes made to a protected VM through the vSphere Client VM Properties are replicated to all copy VMs. Use this procedure to change the default system behavior.

RecoverPoint for VMs replicates the protected VM **version**, **MAC address**, **CPU**, **memory**, **resource reservations**, **network adapter status**, and **network adapter type** to all copy VMs in the consistency group, upon image access.

NOTE: Replication of the SR-IOV NIC type is not supported. If the ESXi at a copy does not support the production VM version, no hardware resources are replicated.

Steps

1. Click **Protection > Protected VMs**.
2. Select a protected VM.
3. Click **PROTECTION POLICY**.
4. Switch **Replicate VM hardware changes** off.
5. Click **UPDATE POLICY**.

Results

Replication of the protected VM hardware changes is enabled or disabled, as defined.

Orchestration

This section describes the RecoverPoint for VMs features for orchestrating virtual machines and VMDKs.

VM start-up sequence

Define the order in which VMs in a consistency group will power on during testing and recovery.

Prerequisites

Install VMware Tools on every production VM. When VMware Tools is installed on a production VM, the VM is considered *powered on* only after its operating system loads. When VMware Tools is not installed on a production VM, the VM is considered *powered on* when it is powered on.

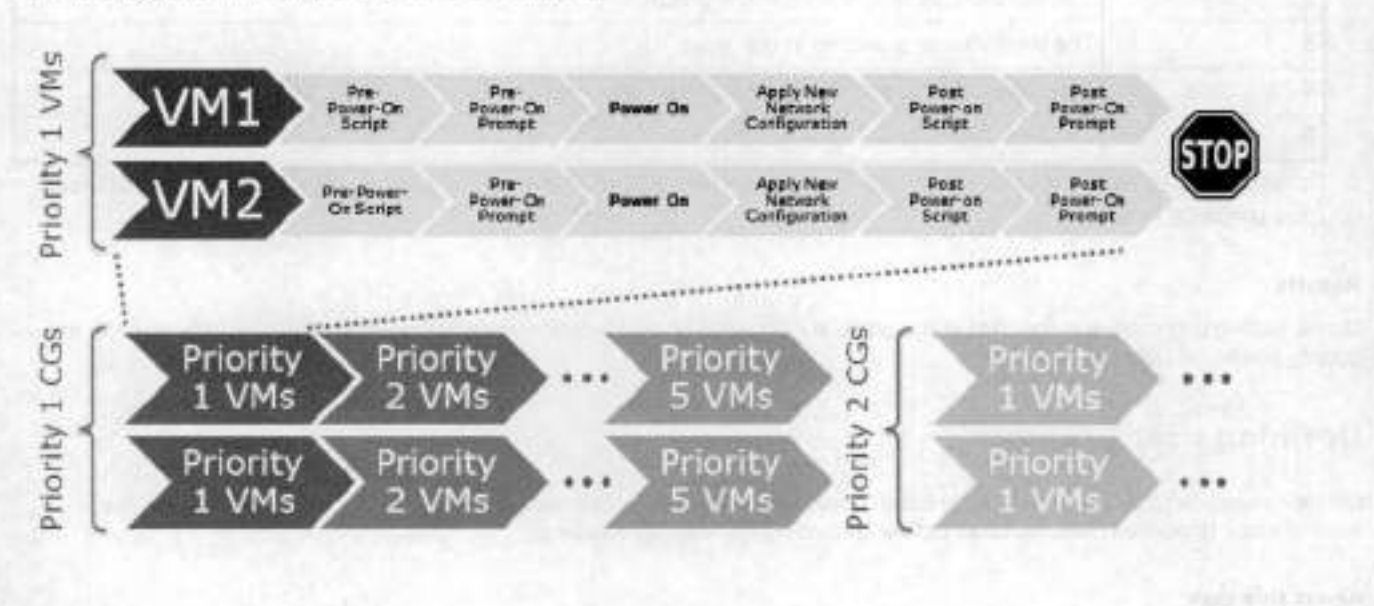
About this task

The VM start-up sequence:

- Is initiated when a copy snapshot is accessed during testing or recovery.
- Moves to the next VM in the start-up sequence only when a VM is *powered on*.

- Enables you to define a VM as **Critical** to ensure that no other VMs power on if the critical VM fails to power on first.
 - Can contain an **Operation**, with one **user script** and one **user prompt** that will run before VM power-on and after VM power-on. Operations run in a strict sequence: **script** > **prompt** > **power-up** > **script** > **prompt**.
- See [Defining user prompts](#) and [Defining user scripts](#) for more information.

The following diagram illustrates the order of sequences:



Steps

1. Click **Protection > Consistency Groups**.
2. Select a consistency group, and click **PROTECTION POLICY**.
The **Group Protection Policy** dialog is displayed.
3. In the **Group Protection Policy** dialog, click **General > VM STARTUP SEQUENCE**.



If the power-on sequence also performs an operation (contains scripts and prompts), a check mark is displayed in the **Operation** column. See [Defining user scripts](#) and [Defining user prompts](#) for more information.

4. Enable **Set same priority for all VMs in group**, or define a **Priority** for each VM in the group.

By default, all priorities are set to **3**.

Priority	Description
1	The first VMs to power on in the group
2	The second VMs to power on in the group
3	The third VMs to power on in the group
4	The fourth VMs to power on in the group
5	The last VMs to power on in the group

5. Optionally, select each VM whose start-up sequence you want to stop if the VM does not power on, and set it to **Critical**.
6. Click **UPDATE POLICY**.

Results

During testing and recovery, the VMs in the group are powered on in the defined order of priority. All of the VMs with the same priority power-on simultaneously.

Defining user prompts

Define a message to display to the administrator that will perform VM copy testing and recovery. User prompts remind the administrator to perform specific tasks before continuing the start-up sequence.

About this task

When defining a VM start-up sequence, you can add a user prompt before power-on and a user prompt after power-on. Administrator's will have to dismiss the prompt before the start-up sequence continues. If a timeout is defined, the prompt is automatically dismissed when the time-out period passes. If no time-out is defined and a prompt is not dismissed, the start-up sequence does not continue until the prompt is dismissed.

Steps

1. Select **Protection > Consistency Groups**, select a consistency group, and click **PROTECTION POLICY**.
2. In the **Protection Policy** dialog, select **General > VM STARTUP SEQUENCE**.
3. Expand a VM, and enable **Prompt user** before or after VM power-on.
4. Type a descriptive name for the prompt.
5. Type the prompt message.
6. Optionally, set the time-out period.

Results

During testing and recovery, administrator's will have to dismiss the prompt before the start-up sequence continues, unless a timeout is defined.

Defining user scripts

Run commands immediately before or after VMs are powered-on during testing and recovery.

Prerequisites

- An external host must be configured. One external host can be defined per vRPA cluster. See *Managing external host registration* for more information.
- An SSH server must be installed on each external host.

About this task

When defining a VM start-up sequence, you can also define the scripts that will be run before or after VMs are powered-on.

- The scripts are run with `ssh` on the external host that you designate.

- Each script has a mandatory time-out period. The recovery flow is blocked until the script runs successfully. A prompt indicates if the script failed.
- You can define one user script before power-on, and one user script after power-on per VM.
- The maximum size of the script name and parameters is 1024 bytes.

Steps

1. Select **Protection > Consistency Groups**, select a consistency group, and click **PROTECTION POLICY**.
2. In the **Protection Policy** dialog, select **General > VM STARTUP SEQUENCE**.
3. Expand a VM, and enable **Run script** before or after VM power-on.
4. Type a descriptive name for the script.
5. Type the script command, including parameters (separated by a space).
6. Set the time-out period (mandatory).
7. Specify the number of retries. If the script does not run within the set time or the script fails, the system retries the script this number of times.

Results

During testing and recovery, these scripts will run before the start-up sequence continues.

Group start-up sequence

Define the order in which VMs of each consistency group in a group set power-on during testing and recovery.

Prerequisites

Install VMware Tools on every production VM. When VMware Tools is installed on a production VM, the VM is considered *powered on* only after its operating system loads. When VMware Tools is not installed on a production VM, the VM is considered *powered on* when it is powered on.

About this task

The group start-up sequence:

- Is initiated when a copy snapshot is accessed during testing or recovery.
- Moves to the next group of VMs in the start-up sequence only when the last group of VMs are all *powered on*.

Steps

1. Click **Protection > Group Sets**.
2. Select a group set.
3. Click **[...] > Group priority**, and define a **Priority** for each consistency group in the group set.

Protection Policy of Group 'cg_TVM2'

Warning: External host has not been defined for one or more vSphere clusters. Select Systems > Orchestration and define an external host for every vSphere cluster of the consistency group.

Set same priority for all VMs in group

Protected VM	Priority	Operation	On/Off	VMware Tools
TVM2	4	N/A	<input checked="" type="checkbox"/>	<input type="checkbox"/>

Warning: VMware Tools is not installed on this virtual machine.

Priority	Description
1	The first group (VMs) to power on in the group set
2	The second (VMs) to power on in the group set
3	The third group (VMs) to power on in the group set
4	The fourth group (VMs) to power on in the group set
5	The last group (VMs) to power on in the group set

4. Click **SAVE**.

Results

During testing and recovery, group VMs are powered on in the defined order of priority. All group VMs with the same priority power-on simultaneously.

Create a group set

Add a group set to RecoverPoint for VMs.

About this task

A group set is a collection of consistency groups that you can bookmark, enable, disable, pause and resume replication for, and test and recover as a group. You can also create parallel bookmarks on all groups in the group set, at a frequency that you define. Group sets are useful for consistency groups that are dependent on one another or that must work together as a single unit.

- NOTE:** You cannot enable parallel bookmarking for a group set containing a group that is part of another group set that has parallel bookmarking enabled.

Steps

1. Select **Protection > Group Sets**.
The **Group Sets** screen is displayed.

Group Sets

UPDATE BOOKMARK TEST A COPY RECOVER PRODUCTION FAILOVER ...

Q Search ADD

Group Set	Consistency Groups	Parallel Bookmark Frequency	vRPA Cluster
My_GroupSet2	1	10 minutes	WELL
MyGroupSet	2	No parallel bookmarks	WELL

Consistency Group	Production vRPA Cluster	Copies
MyGroup2	WELL	2
MyGroup	ALLIS	3

2 Consistency groups in "MyGroupSet"

2 Group sets

- Click **ADD**.
The **Add Group Set** dialog is displayed.
- In the **Add Group Set** dialog:

Add Group Set

vRPA Cluster: **WELL**

Group Set:

Parallel Bookmarks: **Parallel Bookmarks**
Interval:

Consistency Groups

Consistency Group	Production vRPA Cluster	Copies
<input checked="" type="checkbox"/> cg_Win2K12_PROD	WELL	2
<input type="checkbox"/> cg_Win2K12_3804	Toronto	2
<input type="checkbox"/> cg_Win2K12_PROD	Toronto	2
<input checked="" type="checkbox"/> cg_Win2K12_3802	WELL	2

0 Selected

- Choose the vRPA cluster that the consistency groups you want to add to the group set are replicating from. All groups in a group set must be replicating from the same vRPA cluster.
 - Enter a descriptive name for the group set.
 - (Optional) To create bookmarks for all consistency groups in the group set at the same interval, enable **Parallel Bookmarks** and set the desired bookmark interval frequency.
 - Select the consistency groups to add to the group set.
- Click **ADD GROUP SET**.

Results

A new group set is created with the specified settings.

Next steps

See Managing group sets for additional group set capabilities.

Re-IP rules

Create Re-IP rules to update the network configuration of one or more copy VMs when testing a copy, failing over, or recovering production.

Prerequisites

- This feature is supported for VMs running Microsoft Windows server versions 8, 10, 2008 R2, 2012, and 2016, Red Hat Linux server versions 6.5 and 7.2, Red Hat Enterprise Linux (RHEL) server version 7.1, and Ubuntu Studio 15.10.
- VMware Tools should be installed on each relevant production VM.
 - For Linux SLES12, automatic network configuration is not supported unless Open VM Tools version 9.4.0.25793 and `deployPkg` has been manually installed. See [VMware KB article 2075048](#) for detailed information about how to install `deployPkg`.
 - For VMs running Open VM Tools versions lower than 9.10, automatic network configuration is not supported unless `deployPkg` has been manually installed. See [VMware KB article 2075048](#) for detailed information about how to install `deployPkg`.
- Read the Copy VM network configuration guidelines.
- Ensure you do not lose your production VM network configuration during failback by also creating re-ip rules for your production VMs.

About this task

You can specify the testing network of one or more VMs at a copy, or of all copy VMs in the system.

NOTE: Re-IP configuration using glue scripts is not available in the vSphere HTML5 plugin. You can configure glue scripts using the Flex plugin; for instructions, see the *RecoverPoint for VMs Flex Plugin Administrator's Guide*.

Re-IP a few copy VMs

Update the network configuration of a small number of VMs at a copy.

Steps

1. In the RecoverPoint for VMs plugin for vSphere Client, select **Protection > Consistency Groups**.
2. Select a consistency group, and click **PROTECTION POLICY**.
The group **Protection Policy** dialog is displayed.
3. Select the tab of the **Production** or a **Copy** and click the **RE-IP RULES** tab.



4. To retrieve the network configuration of all copy VMs at all vRPA clusters in the system, click **Get values from production**.
5. Update the network values of each copy according to the Copy VM network configuration guidelines.
6. Click **UPDATE POLICY**.

Results

The new copy VM network configuration is used when testing a copy, failing over, or recovering production.

Re-IP many copy VMs

Simultaneously update the network configuration of multiple VMs at a copy.

Steps

1. In the RecoverPoint for VMs plugin for vSphere Client, select **Protection > Consistency Groups**.
2. Select a consistency group, and click **PROTECTION POLICY**.
The group **Protection Policy** dialog is displayed.
3. Select the tab of the **Production** or a **Copy** and click the **RE-IP RULES** tab.
4. To retrieve the network configuration of all copy VMs at all vRPA clusters in the system, click **Get values from production**.

Protection Policy of Group 'cg_newVm'

General Production (Patagonia) Local (Patagonia) Remote 1 (Darwin) **Remote 2 (Darwin)**

COPY POLICY LINK POLICY FAILOVER NETWORKS **RE-IP RULES**

IMPORT EXPORT

Status	Virtual Machine	VMware Tools
▼ ● Not Configured	rp.newVm.copy.copy.shadow	✔ installed Get values from production

5. To facilitate populating JSON with the production values, click **UPDATE POLICY**.
6. Repeat steps 1-3.
7. To save the current network configuration of all virtual machines at the selected copy to a local JSON file, click **Export**.
8. Open the JSON file, and modify the network configuration of relevant virtual machines according to the Copy VM network configuration guidelines.
9. To apply the new network configuration, click **Import** and select the modified JSON file.
10. Click **UPDATE POLICY**.

Results

The new network configuration is used when testing a copy, failing over, or recovering production.

Re-IP all copy VMs in the system

Simultaneously update the network configuration of all VMs of all copies in the system.

Steps

1. In the RecoverPoint for VMs plugin for vSphere Client, select **System > Orchestration** and use the buttons in the **RE-IP Copy VMs** section to update your copy network settings.

External Host

No items to display

ADD

RE-IP Copy VMs

Export, update and import new RE-IP rules for all copies in the system:

IMPORT

EXPORT

2. To save the current network configuration to a local JSON file, click **Export**.
3. Open the JSON file, and modify the network configuration of relevant copy VMs according to the Copy VM network configuration guidelines.
4. To apply the new network configuration to the system, click **Import** and select the modified JSON file.

Results

The new network configuration is used when testing a copy, failing over, or recovering production.

Failover networks

Automatically associate the VM network adapters (vNICs) of copy VMs with specific port groups upon failover, or during copy testing.

About this task

Failover networks can be configured during or after VM protection. Configured failover networks are made available for selection during testing and failover.

Steps

1. In the RecoverPoint for VMs vSphere Client plugin, click **Protection > Consistency Groups**, select a group, and click **PROTECTION POLICY**.
The group **Protection Policy** dialog is displayed.
2. Select a copy and click **FAILOVER NETWORKS**.

Protection Policy of Group 'cg_Win'



General Production (NASA_Site1) **Remote 1 (NASA_Site2)**

COPY POLICY LINK POLICY **FAILOVER NETWORKS** RE-IP RULES

Q Search

Protected VM	Copy VM												
> KATE_VM_1309	rp.KATE_VM_1309.copy1.shadow												
▼ Win	rp.Win.copy.shadow												
<table border="1"><thead><tr><th>Network Adapter</th><th>Production Network</th><th>Network after Failover</th></tr></thead><tbody><tr><td>Network Adapter 1</td><td>Vlan52_nested</td><td>VMs_LAN_DHCP ▼</td></tr></tbody></table>	Network Adapter	Production Network	Network after Failover	Network Adapter 1	Vlan52_nested	VMs_LAN_DHCP ▼	<table border="1"><thead><tr><th>Network Adapter</th><th>Production Network</th><th>Network after Failover</th></tr></thead><tbody><tr><td>Network Adapter 1</td><td>Vlan52_nested</td><td>VMs_LAN_DHCP ▼</td></tr></tbody></table>	Network Adapter	Production Network	Network after Failover	Network Adapter 1	Vlan52_nested	VMs_LAN_DHCP ▼
Network Adapter	Production Network	Network after Failover											
Network Adapter 1	Vlan52_nested	VMs_LAN_DHCP ▼											
Network Adapter	Production Network	Network after Failover											
Network Adapter 1	Vlan52_nested	VMs_LAN_DHCP ▼											

Q vrpa X

- VRPA_LAN
- VRPA_WAN

CANCEL

UPDATE POLICY

- Expand a VM to display its network adapters, and for each adapter, select the network to be used after failover. Use the search filter to easily identify the required network.
- Click **UPDATE POLICY**.

Results

The failover networks are configured.

Next steps

Select **Preconfigured failover networks** when defining the **Testing Network** for copy testing, and when defining the **Target Network** before falling over.

Recovering VMs

Periodically test copy images. In a disaster, fall over to a copy, or recover production to an earlier point-in-time.

Before recovering VMs, see the *RecoverPoint for Virtual Machines Scale and Performance Guide* and the *RecoverPoint for Virtual Machines Release Notes* for information of how to scale your environment, and the limitations of this solution.

Topics:

- Test a copy
- Failover to a copy
- Recover production from a copy

Test a copy

Test a copy of a consistency group or a group set.

Prerequisites

You may want to add journal volumes to a copy journal to ensure that you have ample space for copy testing. For detailed instructions on how to add journal volumes to a copy, see *Managing group protection policies*.

About this task

From time to time, and especially before you begin recovery, test your copy snapshots to ensure they are suitable for recovery. Then, create a bookmark so that suitable snapshots are easily identifiable during recovery.

Steps

1. Depending on whether you want to test a copy of a consistency group or a group set, select **Protection > Consistency Groups** or **Protection > Group Sets**.
2. Select the group or group set whose copy you want to test, and click **TEST A COPY**.
3. In the **Test a Copy** dialog:
 - If you selected a consistency group, select the **Copy to Test**. The vRPA cluster of the copy is displayed, for easy identification.

Test a Copy



Copy to Test

Remote Copy 1 (NASA_Site2) ▾

Snapshot to Test CHANGE

Latest

Testing Networks CHANGE

Isolated per group

Power on copy VMs during testing

CANCEL

START

- If you selected a group set, select the **vRPA Cluster** containing a copy that you want to test. If there are multiple copies at the specified vRPA cluster, RecoverPoint for VMs automatically selects the copy to test. Consistency groups in the group set that do not have a copy at the specified vRPA cluster will be excluded from the activity.

i **NOTE:** To finish this activity, navigate to the **Recovery Activities** screen, **Consistency Groups** tab.

Test a Copy



vRPA Cluster

NASA_Site1 ▾

Snapshot to Test CHANGE

Latest

Testing Networks CHANGE

Isolated per group

Power on copy VMs during testing

CANCEL

START

4. Select the copy **Snapshot to Test**.

Default is **latest**. When selecting a copy snapshot:

- You may want to start with the **Latest** (default) snapshot that is known to be valid.
- You can search for snapshots by name using the search field.
- You can select snapshots by **Bookmark** or **Point in Time**.

Snapshots

Local Time GMT

Search

Point In Time	Size	Bookmark Name	Consistency
<input type="radio"/> Dec 23, 2019 4:4...	296 B		Crash-Consistent
<input type="radio"/> Dec 23, 2019 4:4...	14.95 KB		Crash-Consistent
<input type="radio"/> Dec 23, 2019 4:4...	163.57 KB	MyBookmark	Crash-Consistent

- Use the **zoom in** icon to display all snapshots between the snapshot whose **zoom in** icon you clicked, and the one before it. You can zoom into a snapshot up to 4 times. After zooming into a snapshot you can zoom out by clicking the originating snapshot timestamp.

Snapshots

< DEC 23, 2019 4:30:14 PM

Local Time GMT

Search

Point In Time	Size	Bookmark Name	Consistency
<input type="radio"/> Dec 23, 2019 4:3...	649 B		Crash-Consistent
<input type="radio"/> Dec 23, 2019 4:3...	216 B		Crash-Consistent

- After a snapshot is selected, you can click **REVERT TO LATEST** to revert to the latest snapshot that includes all writes made to it during testing.

Copy to Test

Remote Copy 1 (NASA_Site2)

Snapshot to Test

May 25, 2020 7:34:11 PM REVERT TO LATEST

5. (Optional) Select the copy **Testing Networks**.

To avoid IP conflicts between the production and copy VMs, best practice is to use a dedicated testing network. Therefore, by default, RecoverPoint for VMs auto-provisions an isolated network for all VMs in the group or group set.

Testing Networks

isolated per group

To avoid IP/MAC address conflicts between the production and copy VMs, use an isolated testing network.

Isolated network per group

Automatically creates an isolated network for each consistency group, per ESX.

Isolated network per ESX

Automatically creates an isolated network for each ESX.

Use pre-configured failover networks

Use a dedicated network

vRPA_LAN


You can also:

- Create an isolated network for each ESX.
- Use pre-configured Failover networks.
- Use a dedicated network.

6. Specify whether or not you want RecoverPoint for VMs to **Power on copy VMs during testing**. Default is **enabled**.
7. Click **START** to access the copy snapshot.

Results

The specified snapshot is accessed and a success message is displayed. You can now start testing the copy image.



- NOTE:** If you selected to test a copy of a group set, the success message identifies the copy that the system selected, at the vRPA cluster that you selected.

Next steps

Click the **MORE INFO** link in the success message to go to the **Recovery Activities** screen, and display the activity progress and options. In the **Recovery Activities** screen:

- If you tested a copy of a consistency group:



The **Consistency Group Recovery Activities** screen is displayed. The **Activity Status** and **Progress** columns indicate the progress of image access. After access is enabled to the copy snapshot, the **Activity Status** column displays **Ready for next action**, and you can:

- NOTE:** The **Image Access Log Capacity** progress bar indicates how long you can access the copy image before the image access log is full and all writes to the copy fail. If you need more time for testing, you can undo writes at the copy, enable direct access to the copy, or add journal volumes as described in Managing group protection policies.

- Click **UNDO WRITES** to undo all writes that were made to the copy snapshot while write access was enabled, without disabling access to the copy volumes.
- Click **ENABLE DIRECT ACCESS** to write directly to the copy storage. Any changes made to the copy storage while directly accessing the copy cannot be automatically undone, because when a snapshot is directly accessed, the journal at the copy is deleted. On the other hand, direct access does not impose a limit to the amount of data that you can write to the copy storage volumes.

CAUTION: When direct access is enabled, replication stops to the copy, and a short initialization is required across all group volumes when direct access is disabled.

- Click **ACTIONS > Start new test** to select another snapshot to test, or to redefine the testing network.
- Click **ACTIONS > Stop activity** to roll back all writes that were made to the copy snapshot while write access was enabled to the snapshot volumes, and disable write access to the snapshot volumes.
- Click **ACTIONS > Promote image: Failover** (5.3.1 or later) to Failover to the copy image that you tested in step 8 without needing to roll back the writes that were made to the copy snapshot while write access to the copy volumes was enabled.

- Click **ACTIONS** > **Promote image: Recover Production** (5.3.1 or later) to recover production from the copy image that you tested in step 8 without needing to roll back the writes that were made to the copy snapshot while write access to the copy volumes was enabled.
- If you tested a copy of a group set:

Recovery Activities

Consistency Groups: Group Set

Q Search

Recovery Activity	vRPA Cluster	Consistency Group	Activity Start	Activity Status	Progress	
Test Copy - group-set	GLSA_SNL	3/2	May 25, 2020 6:17 PM	Ready for next action	100%	ACTIONS Stop activity
Summary: Excluded Groups: 0/2 Processing snapshots: 0/2 Processing on VMs: 6/9 Ready for next action: 3/2		Detailed Status: on				
		Testing Network: isolated per group				
		User Prompts: none				

Items per page: 20

Recovery activity

- Click **OPEN** to display a **Detailed Status** for each group in the group set. The **Detailed Status** screen is displayed.
- NOTE:** Groups in the group set without a copy at the specified vRPA cluster are excluded from the activity.

Detailed Status of 'group-set'

Q Search

Consistency Group	Copy	Snapshot	Status	Progress
cg_Win_286	Local Copy	May 25, 2020 6:17 0...	Ready for next action	100%
cg_Win	Remote Copy	May 25, 2020 6:17 0...	Ready for next action	100%

Items per page: 10

2 Detailed Statuses

CLOSE

- Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while access was enabled to the copy volumes, and disable access to the copy volumes.

NOTE: After finding a suitable snapshot, Create a bookmark to label the snapshot so it is easily identifiable for recovery.

Failover to a copy

Failover to a copy of a consistency group or a group set, and (optionally) fallback to the production. You can failover (and fallback) consistency groups or group sets.

About this task

Failover consists of two stages:

- Testing the copy image
- Failover

Failback consists of the same stages.

Before failover, you will have an opportunity to test your copy snapshots to ensure they are suitable for failover.

In environments containing multiple RecoverPoint for VMs systems, to lessen the load on back-end storage arrays, best practice is to failover the consistency groups of up to seven systems concurrently.

Steps

1. Depending on whether you want to failover a group or a group set, select **Protection > Consistency Groups** or **Protection > Group Sets**.
2. Select the group or group set that you want to failover, and click **FAILOVER**.
3. In the **Test a Copy for Failover** dialog:
 - If you selected a consistency group, select the copy and vRPA cluster that you want to failover to in the **Failover to Copy** field.

Test a Copy for Failover



Failover to Copy

Remote Copy 1 (NASA_Site2) ▾

Failover to Snapshot CHANGE

Latest

Testing Networks CHANGE

Isolated per group

CANCEL

START

- If you selected a group set, select the **vRPA Cluster** containing a copy that you want to failover to. If there are multiple copies at the specified vRPA cluster, RecoverPoint for VMs automatically selects the copy. Consistency groups in the group set that do not have a copy at the specified vRPA cluster will be excluded from the activity.

Test a Copy for Failover



vRPA Cluster

NASA_Site2 ▾

Failover to Snapshot CHANGE

Latest

Testing Networks CHANGE

Isolated per group

CANCEL

START

4. Select the snapshot that you want to failover to by clicking **CHANGE** next to **Failover to Snapshot**. Default is the **Latest** snapshot.

When selecting a copy snapshot:

- You may want to start with the **Latest** (default) snapshot that is known to be valid.
- You can search for snapshots by name using the search field.
- You can select snapshots by **Bookmark** or **Point in Time**.

Snapshots

Local Time GMT Q Search C

<input type="radio"/>	Point In Time	Size	Bookmark Name	↑	Consistency	
<input type="radio"/>	Dec 23, 2019 4:4...	216 B			Crash-Consistent	
<input type="radio"/>	Dec 23, 2019 4:4...	14.95 KB			Crash-Consistent	
<input type="radio"/>	Dec 23, 2019 4:4...	163.57 KB	MyBookmark		Crash-Consistent	

- Use the **zoom in** icon to display all snapshots between the snapshot whose **zoom in** icon you clicked, and the one before it. You can zoom into a snapshot up to 4 times. After zooming into a snapshot you can zoom out by clicking the originating snapshot timestamp.

Snapshots

< DEC 23, 2019 4:30:14 PM Local Time GMT Q Search C

<input type="radio"/>	Point In Time	Size	Bookmark Name	↑	Consistency	
<input type="radio"/>	Dec 23, 2019 4:3...	648 B			Crash-Consistent	
<input type="radio"/>	Dec 23, 2019 4:3...	216 B			Crash-Consistent	

- After a snapshot is selected, you can click **REVERT TO LATEST** to revert to the latest snapshot that includes all writes made to it during testing.

Copy to Test

Remote Copy 1 (NASA_Site2) ▾

Snapshot to Test: CLOSE 

May 25, 2020 7:34:11 PM REVERT TO LATEST

5. (Optional) Select the copy **Testing Networks**.

To avoid IP conflicts between the production and copy VMs, best practice is to use a dedicated testing network. Therefore, by default, RecoverPoint for VMs auto-provisions an isolated network for all VMs in the group or group set.

Testing Networks CLOSE

isolated per group

i To avoid IP/MAC address conflicts between the production and copy VMs, use an isolated testing network.

- Isolated network per group**
Automatically creates an isolated network for each consistency group, per REX.
- Isolated network per ESX**
Automatically creates an isolated network for each ESX.
- Use pre-configured failover networks**
- Use a dedicated network**

vRPA_LAN

You can also:

- Create an isolated network for each ESX.
 - Use pre-configured Failover networks.
 - Use a dedicated network.
6. Click **START** to access the copy snapshot.

The specified snapshot is accessed and a success message is displayed.

✓ Accessed snapshot 'Latest' of Remote Copy 1 at vRPA cluster 'NASA_Site2' MORE INFO ✕

Click **MORE INFO** in the success message to go to the **Recovery Activities** screen.

i **NOTE:** If you selected to test a copy of a group set, the success message identifies the copy that the system selected, at the vRPA cluster that you selected.

7. Test the copy image for failover:

In the **Recovery Activities** screen, wait for the **Activity Status** to show **Ready for next action** and the **Progress** status bar, indicating the state of image access to reach 100%.

Then:

- To select a consistency group for failover, ensure the **Consistency Groups** tab is selected.

The screenshot shows the 'Recovery Activities' interface with the 'Consistency Groups' tab selected. It displays a table with columns for Name, Location, Activity Start, Activity Status, and Progress. The first row shows a 'Test Network' activity with a status of 'Ready for next action' and a progress bar at 100%. Below the table, there are several controls: 'Undo Writes' (ENABLED), 'Start transfer' (disabled), and 'Use Snapshot Store' (disabled). An 'Actions' menu is visible on the right side of the table.

i **NOTE:** By default, replication starts immediately after failover. In RecoverPoint for VMs 5.3.1 and later versions, disable **Start transfer** before failing over to pause replication after failover.

- Click **ACTIONS > Start new test** to select another snapshot to test, or to redefine the testing network.
- Click **ACTIONS > Stop activity** to roll back all writes that were made to the copy snapshot while write access was enabled to the snapshot volumes, and disable write access to the snapshot volumes.
- (Optional) Click **UNDO WRITES** to undo all writes that were made to the copy snapshot while write access was enabled, without disabling access to the copy volumes.

- (Optional) Click **ENABLE DIRECT ACCESS** to write directly to the copy storage. Any changes made to the copy storage while directly accessing the copy cannot be automatically undone, because when a snapshot is directly accessed, the journal at the copy is deleted. On the other hand, direct access does not impose a limit to the amount of data that you can write to the copy storage volumes.

CAUTION: When direct access is enabled, replication stops to the copy, and a short initialization is required across all group volumes when direct access is disabled.

- In **Failover Networks**, you can use the default pre-configured failover networks, by keeping **Use or edit pre-configured failover networks** selected. You can also edit a pre-configured network, or choose to **Use current testing networks**.

NOTE: The **Image Access Log Capacity** progress bar indicates how long you can access the copy image before the image access log is full and all writes to the copy fail. If you need more time for testing, you can undo writes at the copy, enable direct access to the copy, or add journal volumes as described in Managing group protection policies.

- To select a group set for failover, click the **Group Sets** tab.

Recovery Activities

Consistency Groups	VRPA Cluster	Consistency Group	Activity Start	Activity Status	Progress
Failover - 001	Start	001	Jul 01, 2021 10:25:00 PM	Ready for next action	60%

NOTE: By default, replication starts immediately after failover. In RecoverPoint for VMs 5.3.2 and later versions, disable **Start transfer** before failing over to pause replication after failover.

- Click **OPEN** to display the **Detailed Status** of all consistency groups in the group set. After access is enabled to the copy snapshot, the **Status** column of all groups displays **Ready for next action**.

NOTE: Groups in the group set without a copy at the specified vRPA cluster are excluded from the activity.

Detailed Status of 'group-set'

Consistency Groups	Copy	Snapshot	Status	Progress
cg_Win_206	Local Copy	May 26, 2020 2:04:2...	Ready for next action	100%
cg_Win	Standalone	May 26, 2020 2:06:1...	Ready for next action	100%

CLOSE

- o Click **ACTIONS > Stop activity** to roll back all writes that were made to the copy snapshot while access was enabled to the copy volumes, and disable access to the copy volumes.

NOTE: After finding a suitable snapshot, you may want to Create a bookmark to label the snapshot so it is easily identifiable during failover.

B. Failover to the copy.

Click **ACTIONS > Failover**.

Results

- If the selected consistency group or group set has only one copy, failover starts.
 - o The role of the **Production** becomes **Remote/Local Copy**.
 - o The role of the **Remote/Local Copy** becomes **Production**.
 - o The production VM and copy VM change roles, but their names do not change. Therefore, after failover, new production VMs will still be named *YourVMName.copy* and the new copy VMs are still named *YourVMName*.
 - o The production journal becomes the copy journal and the copy journal becomes the production journal. You may want to add journal volumes as described in Managing group protection policies.
 - o The marking information in the production journal is deleted, the copy journal is deleted, and the consistency group undergoes a full sweep.

CAUTION: During the full sweep, data is not transferred synchronously.

- If the consistency group or group set has copies other than the copy to which you are failing over (even if they are disabled or replication to them is paused), a temporary failover begins:
 - o The role of the **Production** changes to **Temporary Production**.
 - o The role of the **Remote/Local Copy** changes to **Temporary Remote/Local Copy**.
 - o The roles of any other (unlinked) copies become **Standalone**.
 - o Replication pauses for the other copies and the direction of replication between the production and the failed-over copy changes.

Next steps

After temporary failover, if your consistency group or group set had more than one copy (even if they are disabled or replication to them is paused), in the **Recovery Activities** screen:

- **Failback to the original production.** Select the recovery activity, click **ACTIONS > Test for failback** and run the above procedure beginning with step 3, substituting "failback" for "failover" throughout.

After failing back to the production, if you added volumes to the production journal after failover, to reset the production journal to its original size (by default, 3 GB) without triggering a full sweep click **Protection > Consistency Groups > PROTECTION POLICY**, select the group's **Production** copy, and click **RESET SIZE** in the **Journal Volumes** section.

- **Set the copy as the new production.** Select the recovery activity and click **ACTIONS > Set as production**. If there are standalone (unlinked) copies, the **Set this Copy as the New Production** dialog is displayed.

In the **Set this Copy as the New Production** dialog for consistency groups:

Set this Copy as the New Production ×

Designate this copy at vRPA cluster 'Darwin' as the new production of consistency group 'cg_newVm'.

Before you can permanently failover to this copy, decide what you want to do with the other copies in this group.

For each other copy in the group, define a new copy protection policy, and whether you want to disable or delete the copy after permanent failover.

Standalone 2 (Patagonia)	Async <input type="radio"/> Sync <input type="radio"/>	Disable <input checked="" type="checkbox"/> Enable <input type="checkbox"/>	
Standalone 2 (Darwin)	Async <input type="radio"/> Sync <input type="radio"/>	Disable <input checked="" type="checkbox"/> Enable <input type="checkbox"/>	

CANCEL

SET AS PRODUCTION

1. Configure each standalone copy for consistency groups (or all standalone copies for group sets).

Standalone copies are not linked to the production, and you must decide how to handle them before failover. By default, RecoverPoint for VMs does not delete copy VMs but it does disable them. You can **Enable** any required standalone copies and select a replication mode (sync or async), or **Delete** them from the consistency group. Deleting a copy does not delete the VMs from storage.

CAUTION: Disabled copy VMs require a full sweep when they are re-enabled.

2. Click **SET AS PRODUCTION** to permanently failover.
 - o The role of the **Production** becomes **Remote/Local Copy**.
 - o The role of the **Remote/Local Copy** becomes **Production**.
 - o The standalone copies are handled as specified.
 - o The production VM and copy VM change roles, but their names do not change. Therefore, after failover, new production VMs will still be named *YourVMName_copy* and the new copy VMs are still named *YourVMName*.
 - o The production journal becomes the copy journal and the copy journal becomes the production journal. The production journal does not contain the copy history, so it is by default, a much smaller journal. Therefore, after failover, when the production becomes the copy, you may want to add journal volumes to the new copy journal to ensure that you have ample space for copy testing. For detailed instructions on how to add journal volumes to a copy journal, see *Managing group protection policies*.
 - o The marking information in the production journal is deleted, the journal of the copy to which you failed over is deleted, and the consistency group undergoes a full sweep.

CAUTION: During the full sweep, data is not transferred synchronously.

Recover production from a copy

Production recovery corrects file or logical corruption by rolling the production back to a previous point-in-time. You can recover production of consistency groups or group sets.

About this task

Before you begin recovery, you should test your copy snapshots to ensure they are suitable for recovery.

Production recovery consists of two stages:

- Testing the copy image
- Recovering the production from the copy image

Steps

1. Depending on whether you want to recover the production VMs of a group or a group set, select **Protection > Consistency Groups** or **Protection > Group Sets**.
2. Select the group or group set whose production VMs you want to recover and click **RECOVER PRODUCTION**.
3. In the **Test a Copy for Production Recovery** dialog:
 - If you selected a consistency group, select the copy and vRPA cluster from which you want to recover production in the **Recover Production From Copy** field.

Test a Copy for Production Recovery



Recover Production From Copy

Remote Copy 1 (Darwin) -

Recover Production From Snapshot CHANGE

Latest

Testing Networks CHANGE

Isolated per group

CANCEL

START

- If you selected a group set, select the **vRPA Cluster** containing a copy from which you want to recover production. If there are multiple copies at the specified vRPA cluster, RecoverPoint for VMs automatically selects the copy. Consistency groups in the group set that do not have a copy at the specified vRPA cluster will be excluded from the activity.

Test a Copy for Production Recovery



vRPA Cluster

Darwin -

Recover Production From Snapshot CHANGE

Latest

Testing Networks CHANGE

Isolated per group

CANCEL

START

4. Select the snapshot from which you want to recover production by clicking **CHANGE** next to **Recover Production From Snapshot**. Default is the **Latest** snapshot.

When selecting a copy snapshot:

- You may want to start with the **Latest** (default) snapshot that is known to be valid.
- You can search for snapshots by name using the search field.

- You can select snapshots by **Bookmark** or **Point in Time**.

Snapshots

Local Time GMT

Search

Point In Time	Size	Bookmark Name	T	Consistency	
<input type="radio"/> Dec 23, 2019 4:4	296 B			Crash-Consistent	
<input type="radio"/> Dec 23, 2019 4:4	14.95 KB			Crash-Consistent	
<input type="radio"/> Dec 23, 2019 4:4	163.57 KB	MyBookmark		Crash-Consistent	

- Use the **zoom in** icon to display all snapshots between the snapshot whose **zoom in** icon you clicked, and the one before it. You can zoom into a snapshot up to 4 times. After zooming into a snapshot you can zoom out by clicking the originating snapshot timestamp.

Snapshots

← DEC 23, 2019 4:30:14 PM

Local Time GMT

Search

Point In Time	Size	Bookmark Name	T	Consistency	
<input type="radio"/> Dec 23, 2019 4:3	648 B			Crash-Consistent	
<input type="radio"/> Dec 23, 2019 4:3	296 B			Crash-Consistent	

- After a snapshot is selected, you can click **REVERT TO LATEST** to revert to the latest snapshot that includes all writes made to it during testing.

Copy to Test

Remote Copy 1 (NASA_Site2)

Snapshot to Test

May 25, 2020 7:34:11 PM REVERT TO LATEST

5. (Optional) Select the copy **Testing Networks**.

To avoid IP conflicts between the production and copy VMs, best practice is to use a dedicated testing network. Therefore, by default, RecoverPoint for VMs auto-provisions an isolated network for all VMs in the group or group set.

Testing Networks

Isolated per group

To avoid IP/MAC address conflicts between the production and copy VMs, use an isolated testing network.

Isolated network per group

Automatically creates an isolated network for each consistency group, per ESX.

Isolated network per ESX

Automatically creates an isolated network for each ESX.

Use pre-configured failover networks


Use a dedicated network

vrpa_lan

You can also:

- Create an isolated network for each ESX.
 - Use pre-configured Failover networks.
 - Use a dedicated network.
6. Click **START** to access the copy snapshot.

The specified snapshot is accessed and a success message is displayed.



Click **MORE INFO** in the success message to go to the **Recovery Activities** screen.

NOTE: If you selected to test a copy of a group set, the success message identifies the copy that the system selected, at the vRPA cluster that you selected.

7. Test the copy image for production recovery:

In the **Recovery Activities** screen, wait for the **Activity Status** to show **Ready for next action** and the **Progress** status bar, indicating the state of image access to reach 100%.

Then:

- To select a consistency group for production recovery, ensure the **Consistency Groups** tab is selected.



- o Click **ACTIONS > Start new test** to select another snapshot to test, or to redefine the testing network.
- o Click **ACTIONS > Stop activity** to roll back all writes that were made to the copy snapshot while write access was enabled to the snapshot volumes, and disable write access to the snapshot volumes.
- o (Optional) Click **UNDO WRITES** to undo all writes that were made to the copy snapshot while write access was enabled, without disabling access to the copy volumes.
- o (Optional) Click **ENABLE DIRECT ACCESS** to write directly to the copy storage. Any changes made to the copy storage while directly accessing the copy cannot be automatically undone, because when a snapshot is directly accessed, the journal at the copy is deleted. On the other hand, direct access does not impose a limit to the amount of data that you can write to the copy storage volumes.

CAUTION: When direct access is enabled, replication stops to the copy, and a short initialization is required across all group volumes when direct access is disabled.

NOTE: The **Image Access Log Capacity** progress bar indicates how long you can access the copy image before the image access log is full and all writes to the copy fail. If you need more time for testing, you can undo writes at the copy, enable direct access to the copy, or add journal volumes as described in *Managing group protection policies*.

- To select a group set for production recovery, click the **Group Sets** tab.

Recovery Activities

Consistency Groups **Group Sets**

The screenshot shows a table with the following columns: Recovery Activity, vRPA Cluster, Consistency Group, Activity Start, Activity Date, and Progress. A single row is visible for 'Recover Production - Actual group set' with a status of 'Ready for next action' and 100% progress. An 'Actions' dropdown menu is open, showing 'Stop activity' and 'Recover production' options. A 'Detailed Status' section is also visible, showing 'Excluded Groups: 0/1', 'Preparing snapshot: 0/1', 'Powering on VMs: 0/1', and 'Ready for next action: 1/1'.

Recovery Activity	vRPA Cluster	Consistency Group	Activity Start	Activity Date	Progress
Recover Production - Actual group set	Default	1/1	Jul 20, 2020 1:30:45 PM	Ready for next action	100%

- Click **OPEN** to display the **Detailed Status** of all consistency groups in the group set. After access is enabled to the copy snapshot, the **Status** column of all groups displays **Ready for next action**.

① **NOTE:** Groups in the group set without a copy at the specified vRPA cluster are excluded from the activity.

Detailed Status of 'group-set'

The screenshot shows a table with the following columns: Consistency Groups, Copy, Snapshot, Status, and Progress. Two rows are visible: 'cg_Win_286' with 'Local Copy' and 'May 26, 2020 2:04:2...', and 'cg_Win' with 'Standalone' and 'May 26, 2020 2:06:1...'. Both have a status of 'Ready for next action' and 100% progress. A 'CLOSE' button is visible at the bottom right.

Consistency Groups	Copy	Snapshot	Status	Progress
cg_Win_286	Local Copy	May 26, 2020 2:04:2...	Ready for next action	100%
cg_Win	Standalone	May 26, 2020 2:06:1...	Ready for next action	100%

- Click **ACTIONS > Stop activity** to roll back all writes that were made to the copy snapshot while access was enabled to the copy volumes, and disable access to the copy volumes.

8. Recover production from the copy.

Click **ACTIONS > Recover production**.

Results

- Data transfer from the production to all copies is paused, and will resume only after production recovery is complete.
- Host access to the recovered production volumes, and the recovering copy volumes is blocked.
- Recovered production volumes are overwritten. Any writes made to the copy during testing are transferred to the production, unless you clicked **UNDO WRITES** in step 7.
- The group undergoes a short initialization process to synchronize the new production data at the copy.

Monitoring recovery activities

Monitor ongoing testing, failover, fallback and production recovery activities of consistency groups and group sets, using the **Dashboard**.

Use The RecoverPoint for VMs Dashboard to monitor your recovery activities. The **Dashboard** provides an overview of all ongoing recovery activities in the system. Clicking the status of a recovery activity in a dashboard widget automatically displays the relevant system screen, displaying only the system components in the clicked status. To manage recovery activities, see *Managing recovery activities*.




Figure 7. Monitoring recovery activities

Managing RecoverPoint for VMs

This section describes how to use the **RecoverPoint for VMs vSphere plugin** to manage the components of the RecoverPoint for VMs, after initial system configuration.


Topics:

- Managing your RecoverPoint for VMs licenses
- Managing the plugin server
- Managing vRPA clusters
- Managing system component registration
- Managing protected VMs
- Managing consistency groups
- Managing group sets
- Managing recovery activities
- Configuring email alerts and reports

Managing your RecoverPoint for VMs licenses

Manage RecoverPoint for VMs licenses.

To monitor your RecoverPoint for VMs license usage, use the The RecoverPoint for VMs Dashboard. RecoverPoint for VMs licensing describes the licensing process.



To license RecoverPoint for VMs:

1. Create your license files.
2. Add license.

Managing the plugin server

RecoverPoint for VMs **plugin server** is supported in vSphere 6.7 UI and later versions. One plugin server is supported per vCenter Server (or multiple linked vCenter Servers).

Prerequisites

- Ensure you have a plugin server installed and registered with the vCenter Server that you are connected to. See the *RecoverPoint for VMs Installation and Deployment Guide* for more information on installing plugin servers.
- Ensure you have consulted the *RecoverPoint for VMs Product Guide* for a more detailed description of the plugin server, its architecture, installation, and functionality when vCenter Servers are linked.

- See Managing linked vCenter Server registration.

In the **System > Administration** screen:

- The **vCenter Servers** tab displays all vCenter Servers registered with the plugin server of the vCenter Server that you are connected to.
- The **vRPA Clusters** tab displays all vRPA clusters that are hosted on all linked vCenter Servers registered with the plugin server of the vCenter Server that you are connected to.

Changing the plugin server certificate

Use this procedure to change the plugin server certificate before the plugin server has been configured using **Deployment Manager**.

About this task

Use this procedure, for instance, if you want to use a certificate that has been signed by your organization's internal certificate authority.

Steps

1. Connect to the plugin server with root permissions.
2. Create a backup of the existing certificate and key files:
`/etc/nginx/ssl/rpcenter.cert`
`/etc/nginx/ssl/rpcenter.key`
3. Disable the firewall on the plugin server.
Run the command `/sbin/SuSEfirewall12 off`
4. Upload the new certificate and key files to `/etc/nginx/ssl`.
5. Rename the new certificate file to `rpcenter.cert` and the new key file to `rpcenter.key`
6. Reboot the plugin server VM.
7. In the **RecoverPoint for VMs Deployer**, click **Configure plugin server** home screen.
Enter the **plugin server IP address** in IPv4 format, confirm the new certificate, and click **Configure**.
For more information, see the "Configure the plugin server" in the *RecoverPoint for VMs Installation and Deployment Guide*.

Results

RecoverPoint for VMs is configured to use the new plugin server certificate.

Next steps

NOTE:

Check that the certificate is the same across all vRPAs of the same cluster before adding the vRPA to the cluster.

Log into vSphere Client from the relevant vCenter Server and check that the RecoverPoint for VMs HTML5 plugin is displayed.

Changing a registered plugin server certificate

Use this procedure to change the plugin server certificate after the plugin server has already been configured using **Deployment Manager**.

About this task

Use this procedure, for instance, if you want to use a certificate that has been signed by your organization's internal certificate authority.

Steps

1. Connect to the plugin server with root permissions.
2. Create a backup of the existing certificate and key files:

- ```

/etc/nginx/ssl/rpcenter.cert
/etc/nginx/ssl/rpcenter.key

```
3. Disable the firewall on the plugin server.  
Run the command `/sbin/SUSEfirewall2 off`
  4. Upload the new certificate and key files to `/etc/nginx/ssl`.
  5. Rename the new certificate file to `rpcenter.cert` and the new key file to `rpcenter.key`.
  6. Power off the plugin server VM.
  7. Unregister the RecoverPoint for VMs HTML5 plugin from the relevant vCenter Server.  
See "Unregistering the plugin from the Managed Object Browser" in the *RecoverPoint for VMs Installation and Deployment Guide*.
  8. Power on the plugin server VM.
  9. Navigate to `https://RPCIP/ui`.
  10. Click **Authorize** and enter the vCenter Server Credentials.
  11. Navigate to **DELETE /vcs/{vc-id}** near the bottom of the Swagger page.
  12. Select **Try it Out**, enter the vCenter Server serial number, and select **Execute**.  
A 204 response is returned.
  13. In the **RecoverPoint for VMs Deployer**, click **Configure plugin server** home screen.  
Enter the **plugin server IP address** in IPv4 format, confirm the new certificate, and click **Configure**.  
For more information, see the "Configure the plugin server" in the *RecoverPoint for VMs Installation and Deployment Guide*.

## Results

RecoverPoint for VMs is configured to use the new plugin server certificate.

## Next steps

### NOTE:

Ensure the certificate is the same across all vRPAs of the same cluster before adding the vRPA to the cluster.

Log into vSphere Client from the relevant vCenter Server and check that the RecoverPoint for VMs HTML5 plugin is displayed.

# Managing vCenter Server registration with plugin server

Manage the registration of linked vCenter Servers with the plugin server.

## About this task

Use the **System > Administration > vCenter Servers** tab to register linked vCenter Servers with the plugin server, update the registration of a vCenter Server with the plugin server, or unregister a vCenter Server from the plugin server.

**NOTE:** The first vCenter Server is registered with the plugin server upon system deployment. See the *RecoverPoint for VMs Installation and Deployment Guide* for more information.



Figure 8. Plugin server vCenter Servers screen

## Steps

1. To register a linked vCenter Server with the plugin server, click **Register Linked vCenter**.
2. To update the registration information of a vCenter Server that is already registered with the RecoverPoint for VMs plugin server (for instance, if the vCenter Server password expired, or if you changed the vCenter Server certificate or credentials), and the plugin server session is still active and you can still use the UI:
  - a. Click the **Edit** icon.
  - b. Re-enter the vCenter Server credentials for the plugin server in the **System > Administration > vCenter Servers** tab, or for each vRPA cluster under **System > Administration > vRPA Clusters** tab.
  - c. Use the Sysmgmt CLI command **update\_vcenter\_server\_registration** to enter the new vCenter Server credentials for each relevant vRPA cluster. See the *RecoverPoint for VMs CLI Reference Guide* for more details.
3. To unregister a vCenter Server from the plugin server, click the **Delete** icon.
4. Alternatively, if the plugin server session has expired and the UI is disconnected:
  - a. Use the Sysmgmt CLI command **update\_vcenter\_server\_registration** to enter the new vCenter Server credentials for each relevant vRPA cluster, see the *RecoverPoint for VMs CLI Reference Guide* for more details.
  - b. Use the **Configure plugin server** option in the RecoverPoint for VMs Deployer for a vRPA cluster to update the plugin server with the new certificate, see the *Configure plugin server* procedure in the *RecoverPoint for VMs Installation and Deployment Guide* for more details.

## Next steps

After unregistering a vCenter Server from the plugin server and after updating a plugin server with a new certificate, log out and log back into vSphere, or wait for the session to be re-established.

# Managing linked vCenter Server registration

If you have vCenter Servers that are linked together using **vCenter Embedded Linked Mode**, you can see all vRPA clusters, protected VMs, copy VMs, and the plugin servers that reside on multiple vCenter Servers in a single vSphere **Inventory** view, and protect and recover VMs that reside on multiple vCenter Servers.

**NOTE:** The first vCenter Server per site is registered during system deployment through the **RecoverPoint for VMs Deployer**. Additional linked and non-linked vCenter Servers are registered through the **RecoverPoint for VMs vSphere plugin**.

RecoverPoint for VMs:

- Automatically registers the first vCenter Server when you install the first vRPA cluster.
- Supports up to 7 vCenter Servers linked together using **vCenter Embedded Linked Mode**.
- Uses one user authentication method for all linked vCenter Servers.

To register a linked system, after deploying the vRPA clusters:

1. Ensure you have used the **RecoverPoint for VMs Deployer > Install Plugin Server** button to install and register a plugin server with every vCenter Server hosting a protected VM, a copy VM, a plugin server or a vRPA cluster. The plugin server installs the **vSphere HTML5 plugin** on every registered vCenter Server. When a vCenter Server is not registered with any plugin server, the **vSphere HTML5 plugin** interface is not available through the vSphere Client, and you cannot operate your RecoverPoint for VMs system. See the *RecoverPoint for VMs Installation and Deployment Guide* for more information on installing the plugin server.
2. Register all linked vCenter Servers and the plugin server through the **RecoverPoint for VMs vSphere plugin**:
  - Ensure you have registered all linked vCenter Servers with each vRPA cluster hosted on the linked vCenter Servers.
  - RecoverPoint for VMs automatically registers remote vRPA clusters when a local vRPA cluster is registered, but does not automatically register the vCenter Server on which the remote vRPA clusters are hosted.
  - When a vCenter Server is not registered with at least one vRPA cluster, the **RecoverPoint for VMs > Protect VMs...** menu is disabled, and you cannot protect VMs from the vSphere **Inventory**.
  - Click **System > Administration > vCenter Servers** and ensure you have registered all linked vCenter Servers with a plugin server.
  - When a vCenter Server is not registered with any plugin server, the **RecoverPoint for VMs > Protect VMs...** menu is disabled, and you cannot protect VMs from the vSphere **Inventory**.
3. Connect to a vCenter Server directly to:
  - Manage the plugin server of the linked vCenter Server.
  - Add a RecoverPoint for VMs license to the linked vCenter Server.

Licence usage information is combined for all linked vCenter Servers.

To display all linked vCenter Servers in a linked RecoverPoint for VMs system, click the plugin server **INSTANCE** at the top of the RecoverPoint for VMs vSphere plugin.

INSTANCE 10.10.10.10:123 ▾

| Plugin Instance | Version | vCenter Server    |         |
|-----------------|---------|-------------------|---------|
| 10.10.10.10:123 | 1.0.0   | vc01.mydomain.com | Protect |
| 10.10.10.11:123 | 1.0.0   | vc02.mydomain.com | Active  |

## Local linked vCenter Servers

If you have local vCenter Servers that are linked together using **vCenter Embedded Linked Mode** and remote non-linked vCenter Servers, you can connect to the local vCenter Servers to display all of the production VMs and vRPA clusters at the production site in one vSphere Client inventory. When connecting to the remote vCenter Server, only the remote vRPA clusters and copy VMs will be displayed in the vSphere Client inventory.

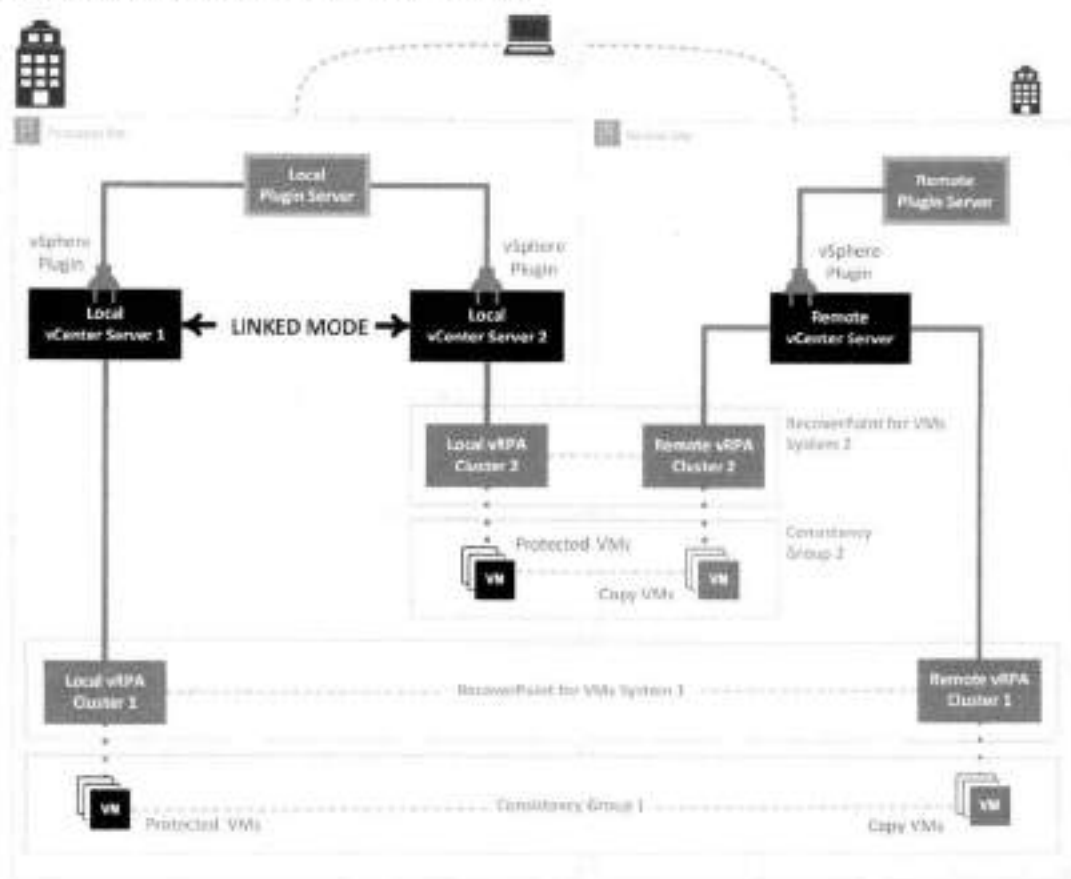



Figure 9. Local linked vCenter example

As illustrated in the Local linked vCenter example, after installing your vRPA clusters:

1. Launch the **RecoverPoint for VMs Deployer** by clicking **System > Administration > vRPA Cluster** and clicking the **Display more vRPA cluster options** icon for a vRPA cluster that is registered to a linked vCenter:



2. Connect to **Local vCenter Server1** or **Local vCenter Server2** and register the vCenter with the **Local Plugin Server**.
3. Connect to the **Remote vCenter Server** and register it with the **Remote Plugin Server**.
4. In the **RecoverPoint for VMs vSphere plugin**:
  - When Managing the plugin server, ensure the **System > Administration > Registered vCenter Servers** table displays linked **Local vCenter Server 2** is registered with the **Local Plugin Server**.
  - When Managing vCenter Server registration with plugin server of **Local vCenter Server1** and **Local vCenter Server2**:
    - Register **Local vCenter Server 1** with **Local vRPA Cluster 1**.
    - Register **Local vCenter Server 2** with **Local vRPA Cluster 2**.
  - Click **INSTANCE** to display both the local and remote plugin servers and all linked and registered vCenter Servers at the production site.

INSTANCE 10.10.10.10:123 ▾

| Plugin Instance | Version | vCenter Server    |         |
|-----------------|---------|-------------------|---------|
| 10.10.10.10:123 | 1.0.0   | vc01.mydomain.com | Protect |
| 10.10.10.11:123 | 1.0.0   | vc02.mydomain.com | Active  |

- When connected to the **Remote vCenter Server**:
  - When Managing vCenter Server registration with plugin server, register **Remote vCenter Server** with **Remote vRPA Cluster 1** and **Remote vRPA Cluster 2**.
  - When Managing the plugin server, register **Remote vCenter Server** with the **Remote Plugin Server**.
  - Click **INSTANCE** to display the remote plugin server and the remote vCenter Server.

| Plugin Instance | Version | vCenter Server                         |
|-----------------|---------|----------------------------------------|
| 10.10.10.10:123 | 1.0.0   | vc01.mydomain.com<br>vc02.mydomain.com |

**NOTE:**

In VMware Cloud Foundation (VCF) multi-tenancy environments (such as VxRail):

- It is best practice to register vRPA clusters with different vCenter Servers (see *Workload Domain or WLD* in VCF documentation).
- All VMs of all WLDs are displayed in the RecoverPoint for VMs plugin, no matter which user is logged in. For example, a user from a tenant on WLD 1 will see and can perform recovery on consistency groups from WLD 2. The RecoverPoint for VMs vSphere plugin does not currently support role-based authentication.
- You can see, manage, protect and recover only the VMs of vRPA clusters that have been directly registered with a vCenter Server.



## Remote linked vCenter Servers

If you have vCenter Servers that are linked together using **vCenter Embedded Linked Mode**, you can display, manage, protect and recover VMs hosted on multiple linked vCenter Servers, across remote sites.

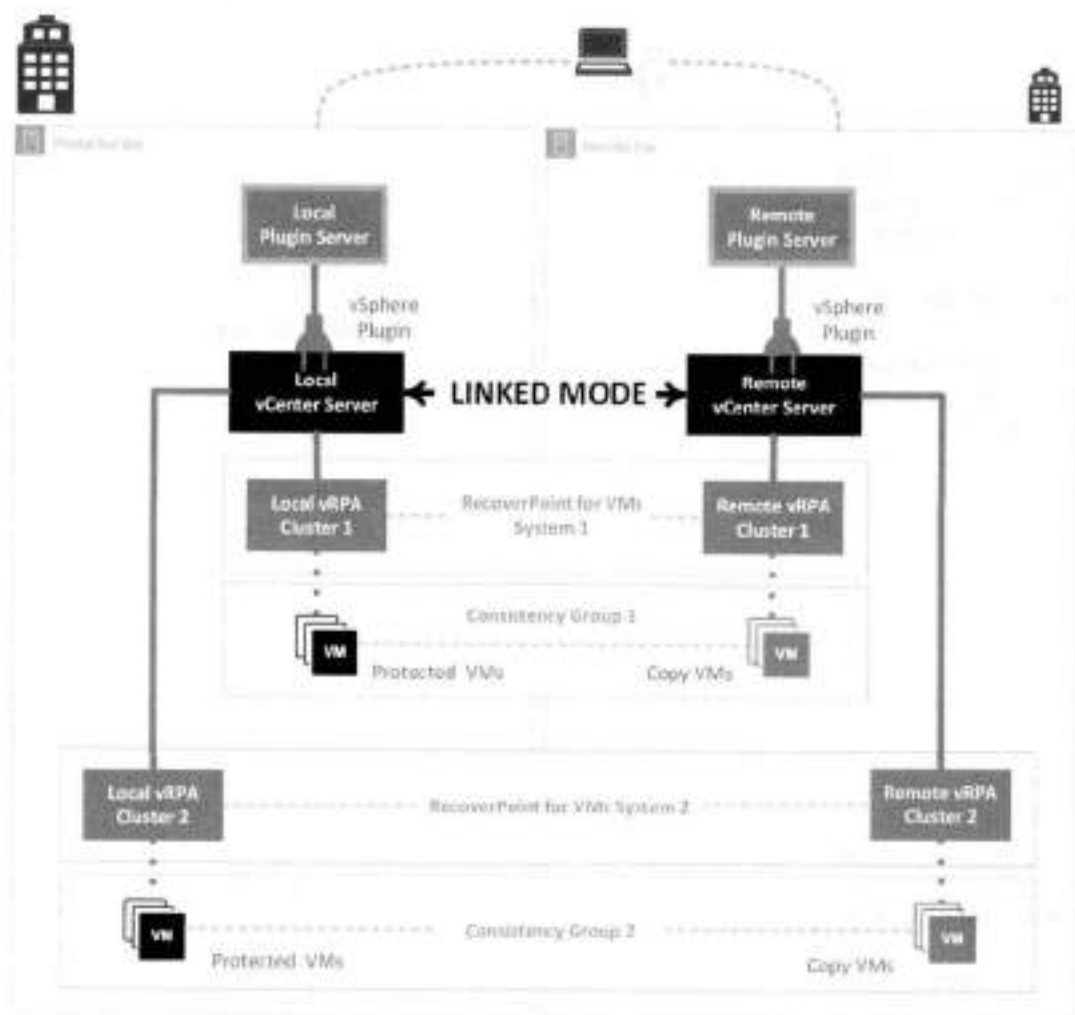


Figure 10. Remote linked vCenter example

As illustrated in the Remote linked vCenter example, after installing your vRPA clusters:

Register the vCenter Servers with the **Local Plugin Server** and **Remote Plugin Server** during system deployment using the **RecoverPoint for VMs Deployer**.

- Register the **Local vCenter Server** with the **Local Plugin Server**.
- Register the **Remote vCenter Server** with the **Remote Plugin Server**.

In the **vSphere plugin**, register linked vCenter (see Managing vCenter Server registration with plugin server).

- Register the **Local vCenter Server** with **Local vRPA Cluster 1** and **Local vRPA Cluster 2**.
- Register the **Remote vCenter Server** with **Remote vRPA Cluster 1** and **Remote vRPA Cluster 2**.

① **NOTE:** To open the **Deployer** from the **vSphere plugin**, you can click **System > Administration > vRPA Clusters**, select a vRPA cluster and click **Display more vRPA cluster options**. See the *RecoverPoint for VMs Installation and Deployment Guide* for more information.

When Managing the plugin server, register the **Remote vCenter Server** with the **Remote Plugin Server**.

After registering your vCenter Servers with the plugin server and the vRPA clusters, you can see and protect VMs of both **System 1** and **System 2** from either vCenter Server.

Click **INSTANCE** to display both plugin server instances and both vCenter Servers are displayed when you are connected to the **RecoverPoint for VMs vSphere plugin** of either vCenter Server.

INSTANCE 10.10.10.10:123

| Plugin Instance | Version | vCenter Server    |        |
|-----------------|---------|-------------------|--------|
| 10.10.10.10:123 | 1.0.0   | vc01.mydomain.com | Active |
| 10.10.10.11:123 | 1.0.0   | vc02.mydomain.com | Active |

## Displaying the plugin server info

To display the plugin server, click the plugin server **INSTANCE** at the top of the **vSphere plugin**.

### About this task

When vCenter Servers are linked, all Managing linked vCenter Server registration that have been registered with the plugin server are displayed. All plugin servers in the linked system are also displayed.

INSTANCE 10.10.10.10:123

| Plugin Instance | Version | vCenter Server    |        |
|-----------------|---------|-------------------|--------|
| 10.10.10.10:123 | 1.0.0   | vc01.mydomain.com | Active |
| 10.10.10.11:123 | 1.0.0   | vc02.mydomain.com | Active |

## Collecting plugin server logs

Collect plugin server logs for support purposes.

### About this task

This procedure collects logs from all vRPAs on all vCenter Servers registered with (or linked to one that is registered with) the plugin server. see [Collecting logs from vRPA clusters](#) to collect logs from specific vRPA clusters.

### Steps

1. Click **System > Administration**

The **RecoverPoint for VMs Plugin Server Administration** screen is displayed.

RecoverPoint for VMs Plugin Server Administration Version: 5.3.2(18)

vRPA Clusters vCenter Servers

Upgrade plugin server  
Collect plugin server logs  
Clear excluded vRPA clusters

| vRPA Cluster | Management IP Address | RecoverPoint for VMs Version | Remote vRPA |
|--------------|-----------------------|------------------------------|-------------|
| Site1        | 10.10.10.10           | 5.3.3P2(m.240)               | Site2       |

Items per page: 20

1 vRPA cluster on vCenter Servers registered with plugin server


2. Click **Actions > Collect plugin server logs**.

The collection of plugin server logs begins and a status indicator is displayed.

Collecting plugin server logs...

## Results

When the collection process is complete, a success message is displayed with the location of the plugin server logs.



## Next steps

Click **Copy**, open a browser window, and paste the copied URL into the browser address bar to retrieve the files.

# Upgrading the plugin server

Upgrading the plugin server upgrades the HTML5 plugin and the API.

## Prerequisites

Download a plugin server upgrade file from the RecoverPoint for VMs product support section of <https://www.dell.com/support>.

## About this task

Plugin server releases are not tied to RecoverPoint for VMs releases. Upgrade packages can upgrade all services running on the plugin server. When a plugin server is being upgraded, the vSphere HTML5 plugin is not functional until upgrade is complete.

## Steps

1. Click **System > Administration**

The **RecoverPoint for VMs Plugin Server Administration** screen is displayed.



2. Click **Actions > Upgrade plugin server**

3. Select the plugin server upgrade file that you downloaded from <https://www.dell.com/support>, and click **OK**.

## Results

During upgrade the RecoverPoint for VMs HTML5 plugin cannot be operated.

## Next steps

Wait for upgrade to complete to operate your RecoverPoint for VMs system.

# Managing vRPA clusters

Manage the vRPA clusters on all vCenter Servers that are registered with (or Linked to a vCenter Server registered with) the plugin server.

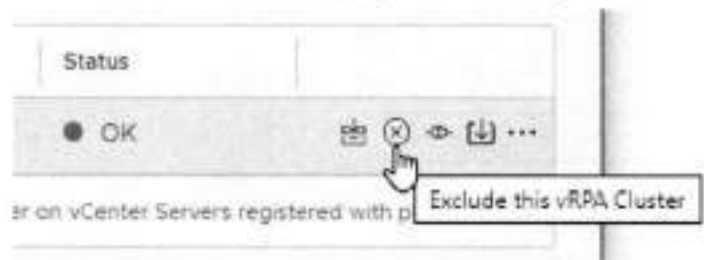
Monitor the state of your **vRPA clusters** in the system in the RecoverPoint for VMs Dashboard.



Use the **System > Administration > vRPA Clusters** screen to collect logs from a vRPA cluster, exclude a vRPA cluster from plugin server management, connect a vRPA cluster to another vRPA cluster, upgrade a vRPA cluster, or display more vRPA cluster options.




- To exclude a vRPA cluster from plugin server management, click the **exclude icon**.



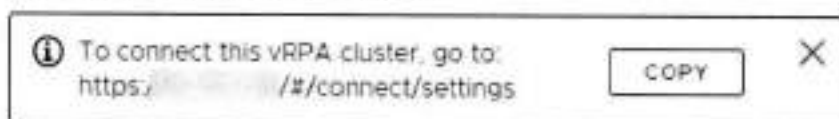
To include all excluded vRPA clusters, click **ACTIONS > Clear excluded vRPA clusters**.



- To connect a vRPA cluster to another vRPA cluster, click the **Connect this vRPA cluster** icon.



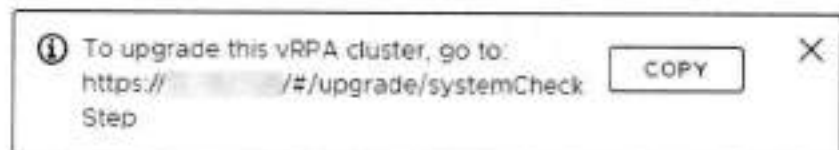
The following system message is displayed:




- To upgrade a vRPA cluster, click the **Upgrade this vRPA cluster** icon.




The following system message is displayed:



- To display more vRPA cluster options click the **ellipses icon [...]**.

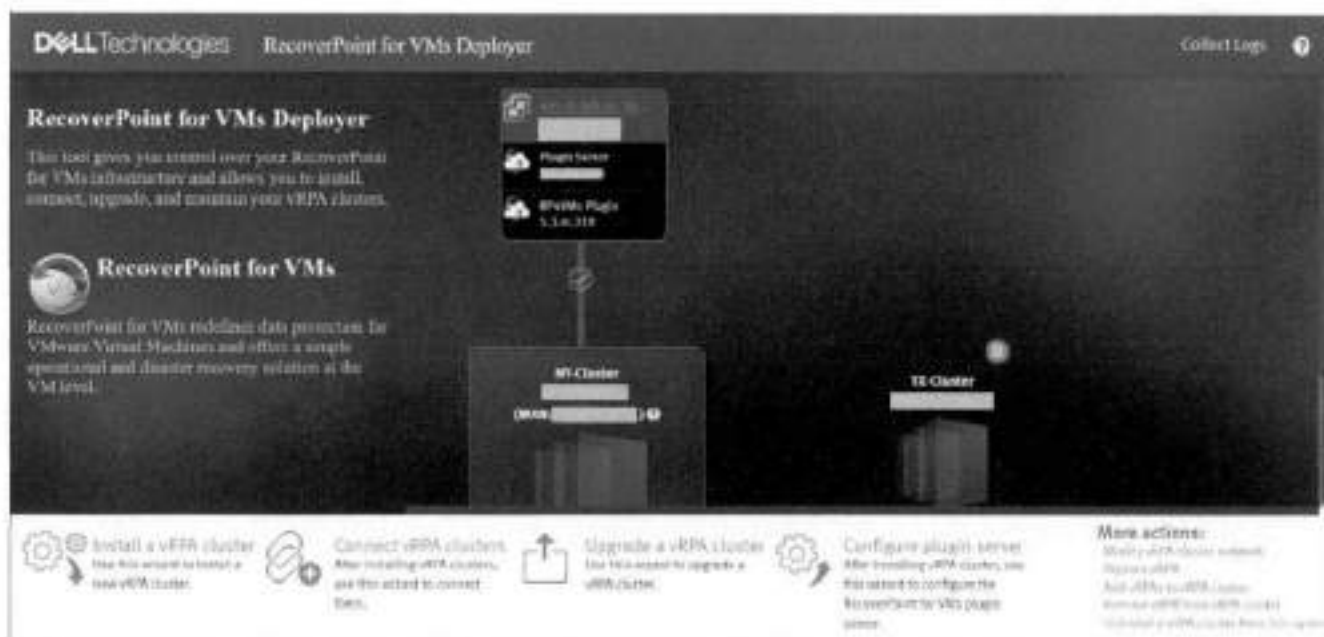


The following system message is displayed:



In the system message that is displayed, click **Copy**, open a browser window, and paste the copied URL into the browser address bar.

The **RecoverPoint for VMs Deployer** is displayed.



In the **RecoverPoint for VMs Deployer** home screen, to complete:

- Click **Collect Logs**. To get the logs from the vRPA, click **Collect Logs** in the **Collect vRPA Logs** dialog box.
- Click **Connect vRPA clusters** and follow the onscreen instructions.
- Click **Upgrade a vRPA cluster** and follow the onscreen instructions.
- Other vRPA cluster options, select the required option under **More actions**.

For detailed information on how to perform all vRPA cluster actions, see the *RecoverPoint for VMs Installation and Deployment Guide*.

## Managing system component registration

This section describes how to manage the registration (and health) of the components of your RecoverPoint for VMs system, after the system has already been configured.

### About this task

After initial system configuration, manage system component registration using the options in the **System** menu.



For a detailed description of how to initially configure the RecoverPoint for VMs system, see [Before you begin](#).

## Register vCenter Server to vRPA cluster

Use this procedure to register a vCenter Server to a vRPA cluster.

### Prerequisites

- All vCenter Servers that manage production VMs and copy VMs must be registered at the relevant vRPA cluster before you protect VMs.
- It is recommended to configure the vCenter Server to require a certificate, because once RecoverPoint has read the certificate, it does not need further access to the location.

For more information about the location of the security certificate, refer to VMware documentation at [www.vmware.com](http://www.vmware.com).

### Steps

1. Click **System** > **vCenter Server**, and select the vRPA cluster to which you want to register a vCenter Server.



2. Click **ADD**.

## Register vCenter Server



### vCenter Server

Enter IP address or host name

### Port

443

### Username

Enter username

### Password

Enter username password

CANCEL

REGISTER

3. Enter the IP and credentials of the vCenter Server to be registered.
4. Click **REGISTER**.

### Results

The specified vCenter Server is registered at the specified vRPA cluster. All ESX clusters hosted by the vCenter Server are automatically registered with the specified vRPA cluster, a splitter is installed on all ESXi hosts in the cluster, and replication is temporarily paused for all relevant VMs while the splitter is being installed.

### Next steps

Ensure that a plugin server is installed on the newly registered vCenter Server. For a linked vCenter Server, see [Managing vCenter Server registration with plugin server](#).

## Managing ESX cluster registration

Registers the ESX cluster of a production VM or copy VM, at a vRPA cluster.

### About this task

By default, ESX clusters are automatically registered in RecoverPoint for VMs during VM protection and copy addition. Use this procedure to register ESX clusters in the rare case that the system cannot automatically register an ESX cluster.

### Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **System > ESX Clusters**.

### Registered ESX Clusters

vRPA Cluster

● Site1\_Cluster

Search

ADD



| ESX Cluster | vCenter Server | Splitter Version | Status |
|-------------|----------------|------------------|--------|
| Site1       | 192.168.1.100  | 5.2.2.0.0.m128   | ● OK   |

Items per page: 20

1 Registered ESX cluster with 'Site1\_Cluster'

2. If you are replicating remotely, select the vRPA cluster at which you want to register ESX clusters.
3. Click **Add**.
4. In the **Register ESX Cluster** dialog box:



- a. Select the ESX cluster that you want to register.
- b. Click **REGISTER**.

### Results

The specified ESX cluster is registered at the specified vRPA cluster.

**NOTE:** When an ESX cluster of an unregistered vCenter Server is registered with a vRPA cluster, a splitter is installed on all ESXs in the cluster, and replication is temporarily paused for all relevant VMs while the splitter is being installed.

### Next steps

Registered ESX clusters can be (1) updated or (2) deleted. To validate ESX cluster registration, see the *RecoverPoint for VMs Flex Plugin Administrator's Guide*.

## Managing journal datastore registration

Register the datastores that are to contain the history of the data that you want to protect, at each vRPA cluster. Up to 15 shared datastores of ESX clusters running vRPAs are automatically registered in RecoverPoint for VMs. Use this procedure to register a datastore in the rare case that a datastore that you need is not automatically registered.

### About this task

**NOTE:** When you protect a consistency group, the **Protect VMs Wizard** will automatically select a datastore from the list of registered datastores, unless you specify a specific registered datastore to use. RecoverPoint for VMs will attempt to create the journal on the selected datastore. If for any reason journal creation fails, the system will attempt to create the journal on a different registered datastore.

### Steps

1. In the **RecoverPoint for VMs vSphere plug-in**, select **System > Datastores**.

Registered Journal Datastores      vRPA Cluster: **Site1\_Cluster**

Automatically register up to 15 shared datastores            **ADD**

| Datastore                   | Total Size | Estimated Free Space |
|-----------------------------|------------|----------------------|
| DEV_RPVE34_Site1_Hosts      | 100.75 GB  | 123.79 GB            |
| DEV_RPVE34_Site1_Journals   | 199.75 GB  | 174.45 GB            |
| DEV_RPVE34_Site1_Repository | 10.75 GB   | 13.32 GB             |
| DEV_RPVE34_Site1_LUNs       | 499.75 GB  | 187.24 GB            |
| DEV_RPVE34_Site1_vRPAs      | 100.75 GB  | 134.59 GB            |

Items per page: 20      5 Registered datastores with 'Site1\_Cluster'

2. If you are replicating remotely, select the vRPA cluster at which you want to register datastores, and click **Add...**. The **Register Datastore** dialog box is displayed.
3. In the **Register Datastore** dialog box:
  - a. Select one or more datastores to register.
  - b. Click **REGISTER**.

### Results

The datastore is registered at the specified vRPA cluster.

### Next steps

Registered journal datastores can be deleted.

**NOTE:** A datastore with a scratch partition on its root path must not be used to host journals.

## Managing external host registration

Defines the external host on which user scripts are run during virtual machine start-up sequences.

### Prerequisites

- SSH must be installed on the external host.
- Only one external host can be configured per vRPA cluster.
- Define the external host before defining virtual machine start-up scripts in a virtual machine startup Sequence. For information on how to define start-up scripts, see VM start-up sequence.

### Steps

1. In the vSphere HTML5 plugin, select **System > Orchestration**, and select the vRPA cluster for which you want to define an external host.
2. Click **ADD** under the **External Host** widget.
3. In the **Register External Host** dialog box, type the **Name**, **IP**, **User**, and **Password** of the external host for the selected vRPA cluster.
4. Optionally:
  - To verify connectivity with the external host, click **Check Connectivity**.
  - To unregister the external host from the specified vRPA cluster, click **Remove**.

## Managing protected VMs

This section describes how to manage the protection of protected VMs.

### About this task

After initial protection, use the RecoverPoint for VMs vSphere HTML5 plugin to manage VMs either through **Protection > Protected VMs** or **Protection > Consistency Groups**. For a detailed description of how to protect VMs, see Protecting VMs.

## Managing VM protection policies

Update the protection policies of protected virtual machines.

### Steps

1. In the RecoverPoint for VMs plugin for vSphere Client, select **Protection > Protected VMs**.
2. Select a virtual machine, and click **PROTECTION POLICY**.  
The VM **Protection Policy** dialog is displayed.

## Protection Policy of VM 'Windows'



Disk Provisioning: Same as source

- Replicate VM hardware changes
- Replicate MAC addresses to the local copy
- Automatically protect newly added VMDKs

Protected VMDKs

| <input checked="" type="checkbox"/> | Protected VMDK | Path      | Size   |
|-------------------------------------|----------------|-----------|--------|
| <input checked="" type="checkbox"/> | Hard disk 1    | SCSI(0:0) | 20 GB  |
| <input checked="" type="checkbox"/> | Hard disk 2    | SCSI(0:1) | 200 MB |

CANCEL

UPDATE POLICY

### 3. Update the VM protection policies:

- **Disk provisioning:** Default is **Same as source**. Defines the way in which the copy VMDKs are to be provisioned; **Same as source**, **Thick provision lazy zeroed**, **Thick provision eager zeroed** or **Thin provision**.
- **Replicate VM hardware changes:** Default is **Enabled**. Automatically replicates the hardware settings of all production virtual machines to their copy VMs whenever an image is accessed on the copy VMs. When enabled, RecoverPoint for VMs replicates the virtual machine version, MAC address, CPU, memory, resource reservations, and network adapter status and type. Replication of SR-IOV Passthrough Adapter is not supported. If the ESX at a copy does not support the production VM version, no hardware resources are replicated.
- **Replicate MAC addresses to local copy:** Default is **Disabled**. If two remote copies of the same production VM are on the same vCenter and in the same network, you cannot power on both copy VMs simultaneously, as they will both have the same MAC address. Therefore, by default, the MAC address of remote copy VMs network adapters (NICs) on a different vCenter than their production VMs are replicated to the copy. However:
  - When **Replicate VM hardware changes** is disabled, MAC address replication to the remote copies is also disabled.
  - To avoid IP conflicts, by default, the MAC addresses are not replicated to the local copy VMs on the same vCenter as their production VMs. If a local copy VM is not on the same network and ESX as its production VM, enable **Replicate MAC addresses to local copy** to replicate the MAC addresses.
- **Automatically protect newly added VMDKs:** Default is **Enabled**. Automatically includes any VMDKs that are added to a VM, after it is already protected.
- **Protected VMDKs:** Displays the number of VMDKs that will be replicated, their **Path**, and their total size. Clear a VMDK check box to exclude the VMDK from replication.

### 4. Click **UPDATE POLICY**.

#### Results

The VM protection policies are updated.

## Stop protecting a VM

Unprotect a VM to stop replication and remove it from its consistency group.

#### Steps

1. In the RecoverPoint for VMs vSphere HTML5 plugin to manage VMs either through **Protection > Protected VMs**
2. Select the production VM that you want to stop protecting.
3. Click **UNPROTECT**

Alternatively, to unprotect all VMs in a consistency group, select **Protection > Consistency Groups**, select the consistency group, and from the [...] menu, select **Unprotect**.

## Results

Replication stops and the virtual machine is removed from its consistency group. The copy VM is not automatically deleted. If there are no other virtual machines in the consistency group, the consistency group is removed.

# Managing consistency groups

This section describes how to manage existing consistency groups in the RecoverPoint for VMs system.

## About this task

After initial protection, use the RecoverPoint for VMs vSphere HTML5 plugin to manage consistency groups through the **Protection > Consistency Groups** screen.

## Managing group protection policies

Update the protection policies of consistency groups and their copies.

### Steps

1. In the RecoverPoint for VMs plugin for **vSphere Client (HTML5)**, select **Protection > Consistency Groups**.
2. Select a consistency group, and click **PROTECTION POLICY**.  
The group **Protection Policy** dialog is displayed.



3. Update the group policies:

- a. Select **General > GROUP POLICY**.

- **Consistency Group:** The name of the consistency group in the RecoverPoint for VMs system. Default is `cg_<vmname>`.
- **Primary vRPA:** The vRPA that you prefer to replicate the consistency group. When the primary vRPA is not available, the consistency group will switch to another vRPA in the vRPA cluster. When the primary vRPA becomes available, the consistency group will switch back to it.

**NOTE:** If your vRPA cluster is the only vRPA cluster in the system, it is a single point of failure in cases of disaster. Consider adding additional vRPAs to the vRPA cluster to ensure high availability.

- **Bandwidth Priority:** Only relevant for remote replication when two or more consistency groups are using the same **Primary vRPA**. Default is **Normal**. Select the bandwidth priority to assign to this consistency group. The priority determines the amount of bandwidth allocated to this group in relation to all other groups using the same primary vRPA.

b. To define the order in which VMs in a consistency group will power on during testing and recovery, select **General > VM STARTUP SEQUENCE**, and refer to VM start-up sequence.

4. Update the group production policies:

**NOTE:** The production policies are only applied after failover.

Select the **Production** tab and click **PRODUCTION POLICY** and **RE-IP RULES** to define the production policies to be applied after failover (when the production becomes a copy). Refer to the next step for a detailed description of the production policies.

5. Update the group copy policies:

a. Select a copy and click the **COPY POLICY** tab to update the copy protection settings:

**Required retention policy**  
7 days

**Consolidate RecoverPoint for VMs snapshots**  
Do not consolidate snapshots for at least:  
2 days

Consolidate snapshots older than 2 days to one snapshot per day for:  
5 days  indefinitely

Consolidate snapshots older than 1 weeks to one snapshot per week for:  
8 weeks  indefinitely

**NOTE:** Snapshots older than 2 months, 3 days will be consolidated to one snapshot per month.

- **Required retention policy:** Default is **disabled**. Defines how far in time the copy image can be rolled back. Enable to define a retention policy in **minutes**, **hours**, **days**, **weeks**, or **months**. An alert is displayed if the copy image cannot be rolled back according to the required retention policy.
- **Consolidate RecoverPoint for VMs snapshots:** Default is **disabled**. Automatic snapshot consolidation cannot be enabled for a group that is part of a group set.
  - **Do not consolidate any snapshots for at least:** Default is **2 days**. Can be defined in **hours**, **days**, **weeks**, or **months**. Defines the period during which snapshot data is not to be consolidated. If no daily or weekly consolidations are specified, the remaining snapshots are consolidated monthly.
  - **Consolidate snapshots that are older than x to one snapshot per day for y days:** Default is **5 days**. Snapshots are consolidated every 24 hours. Select **Indefinitely** to consolidate all subsequent snapshots in 24-hour intervals.
    - If **Indefinitely** is not selected, and no weekly consolidations are specified, the remaining snapshots are consolidated monthly.
    - If **Indefinitely** is selected, weekly and monthly consolidations are disabled, and the remaining snapshots are consolidated daily.
  - **Consolidate snapshots that are older than x to one snapshot per week for y weeks:** Default is **4 weeks**. Snapshots are consolidated every 7 days. Select **Indefinitely** to consolidate all subsequent snapshots in seven-day intervals.
    - If **Indefinitely** is not selected, the remaining snapshots are consolidated monthly.
    - If **Indefinitely** is selected, monthly consolidations are disabled, and the remaining snapshots are consolidated weekly.
- **Journal Volumes:** Displays the size and datastore of each journal volume, and allows you to add or remove journal volumes.

## Journal Volumes

ADD

| Journal Volume      | T | Volume Size | Datastore               |
|---------------------|---|-------------|-------------------------|
| IDFilter_JVOL_00003 |   | 10 GB       | New_DS_Site2_XIO_X1_A07 |

1 Journal volumes

- o (Production journal only) This option is only displayed if more than one journal volume was added to the production journal. Click **RESET SIZE** if journal volumes were added to the production journal after a temporary failover. After falling back to production, use this button to reset the production journal to its original size (by default, **3GB**) without triggering a full sweep.
- o Click the **delete icon** to remove a journal volume. The last journal volume at a copy cannot be deleted.

**CAUTION:** Removing a journal volume causes a full sweep on all VMs in the consistency group. The full sweep duration depends on the size of the data being replicated, network resources, and storage performance.

Click **ADD** to add volumes to the copy journal. The **Add Journal Volume** dialog is displayed.

- o **NOTE:** You can safely click **ADD** now, as the default settings provide a sensible configuration for most systems.

### Add Journal Volume

#### Journal Size

10 GB

#### Select Registered Datastore

Manually  Automatically

CANCEL

ADD

- **Journal Size:** Default is **10GB** for the copy journals and **3GB** for the production journal. The default size of the production journal is smaller, and in the vast majority of cases, will not require any additional volumes. The larger a copy journal, the more history can be saved.
- **Select Registered Datastore:** Default is **Automatically**. By default, RecoverPoint for VMs automatically registers up to 15 datastores for the production and copy journals and automatically selects the datastore with the most free space. When set to **Manually**:

#### Select Registered Datastore

Manually  Automatically

### Registered Datastores

REGISTER DATASTORE

| Datastore               | Total Size | Estimated Free Space |
|-------------------------|------------|----------------------|
| New_DS_Site2_XIO_X1_A07 | 1.5 TB     | 1018 GB              |

Items per page: 20

1 Datastore

- **Registered Datastores:** Select a registered datastore or register a new datastore for the journal. By default, RecoverPoint for VMs automatically registers up to 15 datastores for the group journals and automatically selects the datastore with the most free space. RecoverPoint for VMs will attempt to create

a journal volume on the selected datastore. If it cannot, the system will attempt to create the journal volume on another registered datastore.

- If you have more than 15 datastores and would like to register a datastore that is not in the list, click **REGISTER DATASTORE** to register an additional datastore.

- b. Select a copy and click the **RE-IP RULES** tab to create Re-IP rules to update the network configuration of copy VMs during testing and recovery.
- c. (Not relevant for the production) Select a copy and click the **LINK POLICY** tab to update the copy link protection settings:

Sync  Async

Dynamic by latency

Start async replication above  milliseconds

Resume sync replication below  milliseconds

Dynamic by throughput

Start async replication above  MB

Resume sync replication below  MB

RPO

seconds

Bandwidth Reduction

WAN Compression

Deduplication

- **Async or Sync:** Default is **Async**. Defines the way in which data is replicated from the production to the copy. Data can be replicated synchronously (**sync**) or asynchronously (**async**). When **sync** is selected, you can also define the following policies:
  - **Dynamic by latency:** Default is **disabled**. When enabled, RecoverPoint for VMs alternates between synchronous and asynchronous replication modes, as necessary, according to latency conditions.
    - **Start async replication above:** When the specified limit (in **milliseconds**) is reached, RecoverPoint for VMs automatically starts replicating asynchronously.
    - **Resume sync replication below:** When the specified limit (in **milliseconds**) is reached, RecoverPoint goes back to replicating synchronously.
  - **Dynamic by throughput:** Default is **disabled**. When enabled, RecoverPoint for VMs alternates between synchronous and asynchronous replication modes, as necessary, according to throughput conditions.
    - **Start async replication above:** When the specified limit (in **MB**) is reached, RecoverPoint for VMs automatically starts replicating asynchronously.
    - **Resume sync replication below:** When the specified limit (in **MB**) is reached, RecoverPoint goes back to replicating synchronously.
- **RPO:** Defines the maximum lag allowed on a link.
- **WAN Compression:** Default is **enabled**. Only relevant for asynchronous remote replication. To compress data before transferring it to a remote vRPA cluster, select a level of compression. Enabling and disabling compression causes a short pause in transfer and a short initialization. Compression can reduce transfer time significantly, but increases the source vRPA's CPU utilization.
- **Deduplication:** Default is **disabled**, but deduplication can be enabled whenever compression is enabled. Eliminates repetitive data before transferring the data to a remote vRPA cluster. Enabling and disabling deduplication causes a short pause in transfer and a short initialization. Deduplication can reduce transfer time significantly, but increases the source vRPA's CPU utilization.

- d. (Not relevant for the production) Select a copy and click the **FAILOVER NETWORKS** tab to configure Failover networks to automatically associate the VM network adapters (vNICs) of a copy VM with specific port groups upon failover or during copy testing.

6. Click **UPDATE POLICY**.

### Results

The group protection policies are updated.

## Managing group sets

Manage group sets in RecoverPoint for VMs.

### About this task

A group set is a collection of consistency groups that you can bookmark, enable, disable, pause and resume replication for, and test and recover as a group. You can also create parallel bookmarks on all groups in the group set, at a frequency that you define. Group sets are useful for consistency groups that are dependent on one another or that must work together as a single unit.

### Steps

1. Select **Protection > Group Sets**.

The **Group Sets** screen is displayed.

Group Sets

UPDATE | BOOKMARK | TEST A COPY | RECOVER PRODUCTION | FAILOVER | ...

Search: \_\_\_\_\_ ADD

| Group Set    | Consistency Groups | Parallel Bookmark Frequency | vRPA Cluster |
|--------------|--------------------|-----------------------------|--------------|
| My_GroupSet2 | 1                  | 10 minutes                  | WELL         |
| MyGroupSet   | 2                  | No parallel bookmarks       | WELL         |

| Consistency Group | Production vRPA Cluster | Copies |
|-------------------|-------------------------|--------|
| MyGroup2          | WELL                    | 2      |
| MyGroup           | ALLIS                   | 3      |

2 consistency groups in MyGroupSet

2 Group sets

2. In the **Group Sets** screen:

- Note the number of consistency groups in the group set, the parallel bookmarking status and the vRPA cluster from which all groups in the group set are replicating.
- Expand a group set to display the names of the consistency groups that are currently in the group set and their properties.
- Select a group set, and:
  - Click **ADD** to create another group set. See *Create a group set* for more information.
  - Click **UPDATE** to modify the group set configuration. See *Create a group set* for more information.
  - Click **BOOKMARK** to apply a label or consolidation policy to all copy VMs of all consistency groups in the group set. See *Create a bookmark* for more information.
  - Click **TEST A COPY**, **RECOVER PRODUCTION**, or **FAILOVER** to test a copy, failover, or recover production of all consistency groups in the group set. See *Recovering VMs* for more information.
  - Click the more commands [...] button to display additional group set commands:
    - **Group priority**: Set the power-on priority of all consistency groups in the group set. See *Group start-up sequence* for more information.



- **Disable groups** or **Enable groups**: Disables or enables all copies of all consistency groups in the group set.
- ⚠ **CAUTION: Disabling a group stops all copy activities, and deletes all copy journals. A full sweep is required when transfer resumes.**
- **Pause transfer** or **Resume transfer**: Pauses or resumes transfer of all copies of all consistency groups in the group set.

## Managing recovery activities

Manage ongoing testing, failover, fallback and production recovery activities using the **Recovery Activities** screen.

Use the **Recovery Activities** screen to:

- Manage testing.
- Manage failover and fallback.
- Manage production recovery.

To manage an ongoing recovery activity of a consistency group, select **Recovery ActivitiesConsistency Groups**.

### Recovery Activities

**Consistency Groups**

Groups Sets

| Recovery Activity      | Copy           |
|------------------------|----------------|
| > Recover Production - | Local Copy ( ) |

Items per page 50

To manage an ongoing recovery activity of a group set, select **Recovery ActivitiesGroup Sets**.

### Recovery Activities

Consistency Groups

**Groups Sets**

| Recovery Activity | vRPA Cluster |
|-------------------|--------------|
| > Test Copy - aut | VxRail_Site3 |


Items per page 50

Use the **The RecoverPoint for VMs Dashboard** to monitor the state of ongoing recovery activities, for both consistency groups and group sets.

## Group Recovery Activities ①



## Group Set Recovery Activities ①



## Manage testing

Manage the ongoing testing of a copy of a consistency group or group set using the **Recovery Activities** screen.

### ① NOTE:

When testing a copy, wait for the **Activity Status** to show **Ready for next action** and the **Progress status bar** to reach 100%, indicating the specified snapshot image has been accessed.

- If you tested a copy of a consistency group:

### Recovery Activities

Consistency Groups ① ② ③ ④ ⑤ ⑥ ⑦ ⑧ ⑨ ⑩ ⑪ ⑫ ⑬ ⑭ ⑮ ⑯ ⑰ ⑱ ⑲ ⑳ ㉑ ㉒ ㉓ ㉔ ㉕ ㉖ ㉗ ㉘ ㉙ ㉚ ㉛ ㉜ ㉝ ㉞ ㉟ ㊱ ㊲ ㊳ ㊴ ㊵ ㊶ ㊷ ㊸ ㊹ ㊺ ㊻ ㊼ ㊽ ㊾ ㊿

| Recovery Activity                                                                                                              | State                   | Snapshot                              | Activity Start        | Activity Status       | Progress |                |
|--------------------------------------------------------------------------------------------------------------------------------|-------------------------|---------------------------------------|-----------------------|-----------------------|----------|----------------|
| Test Copy of STATE_VM_01                                                                                                       | Copy From (STATE_VM_01) | Nov 21, 2020 11:47 AM                 | Nov 21, 2020 11:48 AM | Ready for next action | 100%     | <b>ACTIONS</b> |
| Image Access Log Capacity                                                                                                      |                         | Testing/Network: (connect to network) |                       |                       |          |                |
| <div style="display: flex; justify-content: space-between;"> <span>UNDO WRITES</span> <span>ENABLE DIRECT ACCESS</span> </div> |                         | User: (username)                      |                       |                       |          |                |
| Journal Capacity                                                                                                               |                         |                                       |                       |                       |          |                |

The **Consistency Group Recovery Activities** screen is displayed. The **Activity Status** and **Progress** columns indicate the progress of image access. After access is enabled to the copy snapshot, the **Activity Status** column displays **Ready for next action**, and you can:

① **NOTE:** The **Image Access Log Capacity** progress bar indicates how long you can access the copy image before the image access log is full and all writes to the copy fail. If you need more time for testing, you can undo writes at the copy, enable direct access to the copy, or add journal volumes as described in *Managing group protection policies*.

- Click **UNDO WRITES** to undo all writes that were made to the copy snapshot while write access was enabled, without disabling access to the copy volumes.
- Click **ENABLE DIRECT ACCESS** to write directly to the copy storage. Any changes made to the copy storage while directly accessing the copy cannot be automatically undone, because when a snapshot is directly accessed, the journal at the copy is deleted. On the other hand, direct access does not impose a limit to the amount of data that you can write to the copy storage volumes.

⚠ **CAUTION:** When direct access is enabled, replication stops to the copy, and a short initialization is required across all group volumes when direct access is disabled.

- Click **ACTIONS** > **Start new test** to select another snapshot to test, or to redefine the testing network.
- Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while write access was enabled to the snapshot volumes, and disable write access to the snapshot volumes.

- o Click **ACTIONS > Promote image: Failover** (5.3.1 or later) to jump directly to the Failover stage of failover to the copy image that you just tested (step 8), without needing to roll back the writes that were made to the copy snapshot while write access was enabled.
  - o Click **ACTIONS > Promote image: Recover Production** (5.3.1 or later) to jump directly to the production recovery stage of recovering production from the copy image that you just tested (step 8), without needing to roll back the writes that were made to the copy snapshot while write access to the copy volumes was enabled.
- If you tested a copy of a group set:

Recovery Activities

Consistency Groups   **Group Sets**

Search

| Recovery Activity          | VRPA Queue | Consistency Group                             | Activity Start          | Activity Status       | Progress |                                                           |
|----------------------------|------------|-----------------------------------------------|-------------------------|-----------------------|----------|-----------------------------------------------------------|
| Test Copy - Group Set      | SACR_2192  | 2/2                                           | May 25, 2020 8:21:52 PM | Ready for next action | 100%     | <a href="#">ACTIVATE</a><br><a href="#">Stop Activity</a> |
| <b>Summary:</b>            |            | <b>Detailed Status:</b> <a href="#">click</a> |                         |                       |          |                                                           |
| Excluded Groups: 0/2       |            | Testing Network: disabled per group           |                         |                       |          |                                                           |
| Preparing snapshot: 0/2    |            | User Draggable Icons                          |                         |                       |          |                                                           |
| Powering on VM: 0/2        |            |                                               |                         |                       |          |                                                           |
| Ready for next action: 2/2 |            |                                               |                         |                       |          |                                                           |

Items per page: 10

Recovery Activity

- o Click **OPEN** to display a **Detailed Status** for each group in the group set. The **Detailed Status** screen is displayed.
- NOTE:** Groups in the group set without a copy at the specified vRPA cluster are excluded from the activity.

### Detailed Status of 'group-set'

Search

| Consistency Groups | Copy        | Snapshot             | Status                | Progress |
|--------------------|-------------|----------------------|-----------------------|----------|
| cg_Win_266         | Local Copy  | May 25, 2020 8:17:00 | Ready for next action | 100%     |
| cg_Win             | Remote Copy | May 25, 2020 8:17:00 | ready for next action | 100%     |

Items per page: 10

2 Detailed Statuses

CLOSE

- o Click **ACTIONS > Stop activity** to roll back all writes that were made to the copy snapshot while access was enabled to the copy volumes, and disable access to the copy volumes.
- If you tested a copy of a consistency group:

## Recovery Activities

| Recovery Activity         | Plan                   | Snapshot                 | Activity Start           | Activity Status       | Progress |                |
|---------------------------|------------------------|--------------------------|--------------------------|-----------------------|----------|----------------|
| Test Copy - cg_2020_09_25 | Test Copy (Auto State) | Nov 25, 2020 11:02:47 AM | Nov 25, 2020 11:02:48 AM | Ready for next action | 100%     | <b>ACTIONS</b> |

Image Access Log Capacity: 100%  
 Journal Capacity: 75%

Testing Network: isolated per group  
 User Throughput: None

Start new test  
 Stop activity  
 Promote image  
 Failover  
 Recover Production

The **Consistency Group Recovery Activities** screen is displayed. The **Activity Status** and **Progress** columns indicate the progress of image access. After access is enabled to the copy snapshot, the **Activity Status** column displays **Ready for next action**, and you can:

**NOTE:** The **Image Access Log Capacity** progress bar indicates how long you can access the copy image before the image access log is full and all writes to the copy fail. If you need more time for testing, you can undo writes at the copy, enable direct access to the copy, or add journal volumes as described in Managing group protection policies.

- Click **UNDO WRITES** to undo all writes that were made to the copy snapshot while write access was enabled, without disabling access to the copy volumes.
- Click **ENABLE DIRECT ACCESS** to write directly to the copy storage. Any changes made to the copy storage while directly accessing the copy cannot be automatically undone, because when a snapshot is directly accessed, the journal at the copy is deleted. On the other hand, direct access does not impose a limit to the amount of data that you can write to the copy storage volumes.

**CAUTION:** When direct access is enabled, replication stops to the copy, and a short initialization is required across all group volumes when direct access is disabled.

- Click **ACTIONS** > **Start new test** to select another snapshot to test, or to redefine the testing network.
- Click **ACTIONS** > **Stop activity** to roll back all writes that were made to the copy snapshot while write access was enabled to the snapshot volumes, and disable write access to the snapshot volumes.
- Click **ACTIONS** > **Promote image: Failover** (5.3.1 or later) to jump directly to the Failover stage of failover to the copy image that you just tested (step 8), without needing to roll back the writes that were made to the copy snapshot while write access was enabled.
- Click **ACTIONS** > **Promote image: Recover Production** (5.3.1 or later) to jump directly to the production recovery stage of recovering production from the copy image that you just tested (step 8), without needing to roll back the writes that were made to the copy snapshot while write access to the copy volumes was enabled.

- If you tested a copy of a group set:

## Recovery Activities

| Recovery Activity    | VRPA Cluster | Consistency Group | Activity Start          | Activity Status       | Progress |                |
|----------------------|--------------|-------------------|-------------------------|-----------------------|----------|----------------|
| Test Copy - 09240101 | VRPA_0102    | CG1               | Nov 25, 2020 9:21:53 AM | Ready for next action | 100%     | <b>ACTIONS</b> |

Summary: Detailed Status: Open  
 Protected Objects: 1/2  
 Preparing snapshot: 0/2  
 Powering on VMs: 0/2  
 Ready for next action: 2/2

Stop activity

- Click **OPEN** to display a **Detailed Status** for each group in the group set. The **Detailed Status** screen is displayed.

**NOTE:** Groups in the group set without a copy at the specified VRPA cluster are excluded from the activity.

## Detailed Status of 'group-set'

| Consistency Groups | Copy        | Snapshot            | Status                | Progress |
|--------------------|-------------|---------------------|-----------------------|----------|
| cg_Wrk_396         | Local Copy  | May 25, 2020 6:17:0 | Ready for next action | 100%     |
| cg_Wrk             | Remote Copy | May 25, 2020 6:17:0 | Ready for next action | 100%     |

Items per page: 30 | 2 Detailed Statuses

CLOSE

- Click **ACTIONS > Stop activity** to roll back all writes that were made to the copy snapshot while access was enabled to the copy volumes, and disable access to the copy volumes.
- If you tested a copy of a consistency group:

### Recovery Activities

| Recovery Activity           | Type                    | Created               | Starting Date            | Activity Status       | Progress | ACTIONS                                                                                                                                                      |
|-----------------------------|-------------------------|-----------------------|--------------------------|-----------------------|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Test Failover: cg_Wrk_VA_01 | Local Image (SMTL) Copy | Nov 25, 2020 10:47 AM | Nov 25, 2020 11:25:00 AM | Ready for next action | 100%     | <ul style="list-style-type: none"> <li>Start new test</li> <li>Stop activity</li> <li>Promote image</li> <li>Failover</li> <li>Recover Production</li> </ul> |

Items per page: 20

The **Consistency Group Recovery Activities** screen is displayed. The **Activity Status** and **Progress** columns indicate the progress of image access. After access is enabled to the copy snapshot, the **Activity Status** column displays **Ready for next action**, and you can:

- NOTE:** The **Image Access Log Capacity** progress bar indicates how long you can access the copy image before the image access log is full and all writes to the copy fail. If you need more time for testing, you can **undo writes** at the copy, **enable direct access** to the copy, or **add journal volumes** as described in *Managing group protection policies*.
- Click **UNDO WRITES** to undo all writes that were made to the copy snapshot while write access was enabled, without disabling access to the copy volumes.
- Click **ENABLE DIRECT ACCESS** to write directly to the copy storage. Any changes made to the copy storage while directly accessing the copy cannot be automatically undone, because when a snapshot is directly accessed, the journal at the copy is deleted. On the other hand, direct access does not impose a limit to the amount of data that you can write to the copy storage volumes.
- CAUTION:** When direct access is enabled, replication stops to the copy, and a short initialization is required across all group volumes when direct access is disabled.
- Click **ACTIONS > Start new test** to select another snapshot to test, or to redefine the testing network.
- Click **ACTIONS > Stop activity** to roll back all writes that were made to the copy snapshot while write access was enabled to the snapshot volumes, and disable write access to the snapshot volumes.
- Click **ACTIONS > Promote image: Failover** (5.3.1 or later) to jump directly to the Failover stage of failover to the copy image that you just tested (step 8), without needing to roll back the writes that were made to the copy snapshot while write access was enabled.
- Click **ACTIONS > Promote image: Recover Production** (5.3.1 or later) to jump directly to the production recovery stage of recovering production from the copy image that you just tested (step 8), without needing to roll back the writes that were made to the copy snapshot while write access to the copy volumes was enabled.

- If you tested a copy of a group set:

## Recovery Activities

Consistency Groups **Group Sets**

Search

| Recovery Activity                                                                                                         | vRPA Cluster | Consistency Groups                                                                       | Activity Start          | Activity Status       | Progress                                                                     |                                                          |
|---------------------------------------------------------------------------------------------------------------------------|--------------|------------------------------------------------------------------------------------------|-------------------------|-----------------------|------------------------------------------------------------------------------|----------------------------------------------------------|
| Test Copy - Group Set                                                                                                     | NSA_V102     | CG                                                                                       | May 25, 2020 6:17:50 PM | Ready for next action | <div style="width: 100%; height: 10px; background-color: black;"></div> 100% | <a href="#">Actions</a><br><a href="#">Stop activity</a> |
| <b>Summary:</b><br>Excluded Groups: 0/2<br>Preserving snapshot: 0/2<br>Powering on VMs: 0/2<br>Ready for next action: 2/2 |              | <b>Detailed Status:</b> OPEN<br>Testing network: isolated per group<br>User Prompt: None |                         |                       |                                                                              |                                                          |

Items per page: 20

- o Click **OPEN** to display a **Detailed Status** for each group in the group set. The **Detailed Status** screen is displayed.
- NOTE:** Groups in the group set without a copy at the specified vRPA cluster are excluded from the activity.

### Detailed Status of 'group-set'

Search

| Consistency Groups | Copy        | Snapshot               | Status                | Progress                                                                     |
|--------------------|-------------|------------------------|-----------------------|------------------------------------------------------------------------------|
| cg_Win_286         | Local Copy  | May 25, 2020 6:17:0... | Ready for next action | <div style="width: 100%; height: 10px; background-color: black;"></div> 100% |
| cg_Win             | Remote Copy | May 25, 2020 6:17:0... | Ready for next action | <div style="width: 100%; height: 10px; background-color: black;"></div> 100% |

Items per page: 10

2 Detailed Statuses

**CLOSE**

- o Click **ACTIONS > Stop activity** to roll back all writes that were made to the copy snapshot while access was enabled to the copy volumes, and disable access to the copy volumes.

After finding a suitable snapshot, Create a bookmark to label the snapshot so it is easily identifiable for recovery.

## Manage failover and failback

Manage the ongoing failover or failback of a copy of a consistency group or group set using the **Recovery Activities** screen.

### **NOTE:**

When failing over (or failing back) to a copy, wait for the **Activity Status** to show **Ready for next action** and the **Progress status bar** to reach 100%, indicating the specified snapshot image has been accessed.

When the **Activity Status** is **Ready for next action**:

- To select a consistency group for failover, ensure the **Consistency Groups** tab is selected.

## Recovery Activities

Consistency Groups [Group Sets](#)

The screenshot shows a table with columns: Recovery Activity, Consistency Group, Status, Activity Date, Activity Status, Progress, and Actions. The first row shows a consistency group named 'P10000-10\_0000\_0000\_0000' with a status of 'Ready for next action'. Below the table, there are sections for 'Image Access Log Capacity', 'Journal Capacity', 'Testing Network', 'Failover Networks', and 'Start transfer'.

**NOTE:** By default, replication starts immediately after failover. In RecoverPoint for VMs 5.3.1 and later versions, disable **Start transfer** before failing over to pause replication after failover.

- Click **ACTIONS > Start new test** to select another snapshot to test, or to redefine the testing network.
- Click **ACTIONS > Stop activity** to roll back all writes that were made to the copy snapshot while write access was enabled to the snapshot volumes, and disable write access to the snapshot volumes.
- (Optional) Click **UNDO WRITES** to undo all writes that were made to the copy snapshot while write access was enabled, without disabling access to the copy volumes.
- (Optional) Click **ENABLE DIRECT ACCESS** to write directly to the copy storage. Any changes made to the copy storage while directly accessing the copy cannot be automatically undone, because when a snapshot is directly accessed, the journal at the copy is deleted. On the other hand, direct access does not impose a limit to the amount of data that you can write to the copy storage volumes.

**CAUTION:** When direct access is enabled, replication stops to the copy, and a short initialization is required across all group volumes when direct access is disabled.

- In **Failover Networks**, you can use the default pre-configured failover networks, by keeping **Use or edit pre-configured failover networks** selected. You can also edit a pre-configured network, or choose to **Use current testing networks**.

**NOTE:** The **Image Access Log Capacity** progress bar indicates how long you can access the copy image before the image access log is full and all writes to the copy fail. If you need more time for testing, you can undo writes at the copy, enable direct access to the copy, or add journal volumes as described in Managing group protection policies.

- To select a group set for failover, click the **Group Sets** tab.

## Recovery Activities

Consistency Groups [Group Sets](#)

The screenshot shows a table with columns: Recovery Activity, vRPA Cluster, Consistency Group, Activity Date, Activity Status, Progress, and Actions. The first row shows a group set named 'P10000-100' with a status of 'Ready for next action'. Below the table, there is a 'Summary' section with 'Detailed Status' and 'Testing Network' options, and a 'Start transfer' toggle.

- Click **OPEN** to display the **Detailed Status** of all consistency groups in the group set. After access is enabled to the copy snapshot, the **Status** column of all groups displays **Ready for next action**.

**NOTE:** Groups in the group set without a copy at the specified vRPA cluster are excluded from the activity.

## Detailed Status of 'group-set'

| Consistency Group | Copy       | Snapshot               | Status                | Progress |
|-------------------|------------|------------------------|-----------------------|----------|
| cg_Win_286        | Local Copy | May 26, 2020 1:04:2... | Ready for next action | 100%     |
| cg_Win            | Standalone | May 26, 2020 1:06:1... | Ready for next action | 100%     |

Items per page: 10

2 Detailed Statistics

CLOSE

- Click **ACTIONS > Stop activity** to roll back all writes that were made to the copy snapshot while access was enabled to the copy volumes, and disable access to the copy volumes.

**NOTE:** After finding a suitable snapshot, you may want to Create a bookmark to label the snapshot so it is easily identifiable during failover.

Click **ACTIONS > Failover** to failover to the copy

- If the selected consistency group or group set has only one copy, failover starts.
  - The role of the **Production** becomes **Remote/Local Copy**.
  - The role of the **Remote/Local Copy** becomes **Production**.
  - The production VM and copy VM change roles, but their names do not change. Therefore, after failover, new production VMs will still be named *YourVMName.copy* and the new copy VMs are still named *YourVMName*.
  - The production journal becomes the copy journal and the copy journal becomes the production journal. You may want to add journal volumes as described in Managing group protection policies.
  - The marking information in the production journal is deleted, the copy journal is deleted, and the consistency group undergoes a full sweep.

**CAUTION:** During the full sweep, data is not transferred synchronously.

- If the consistency group or group set has copies other than the copy to which you are failing over (even if they are disabled or replication to them is paused), a temporary failover begins:
  - The role of the **Production** changes to **Temporary Production**.
  - The role of the **Remote/Local Copy** changes to **Temporary Remote/Local Copy**.
  - The roles of any other (unlinked) copies become **Standalone**.
  - Replication pauses for the other copies and the direction of replication between the production and the failed-over copy changes.

After temporary failover, if your consistency group or group set had more than one copy (even if they are disabled or replication to them is paused), in the **Recovery Activities** screen:

- Failback to the original production.** Select the recovery activity, click **ACTIONS > Test for fallback** and run the above procedure beginning with step 3, substituting "fallback" for "failover" throughout.

After failing back to the production, if you added volumes to the production journal after failover, to reset the production journal to its original size (by default, 3 GB) without triggering a full sweep click **Protection > Consistency Groups > PROTECTION POLICY**, select the group's **Production** copy, and click **RESET SIZE** in the **Journal Volumes** section.

- Set the copy as the new production.** Select the recovery activity and click **ACTIONS > Set as production**. If there are standalone (unlinked) copies, the **Set this Copy as the New Production** dialog is displayed.

In the **Set this Copy as the New Production** dialog for consistency groups:



## Set this Copy as the New Production ×

Designate this copy at vRPA cluster 'Darwin' as the new production of consistency group 'cg\_newVm'.

Before you can permanently failover to this copy, decide what you want to do with the other copies in this group.

For each other copy in the group, define a new copy protection policy, and whether you want to disable or delete the copy after permanent failover.

|                             |                                                        |                                                                             |  |
|-----------------------------|--------------------------------------------------------|-----------------------------------------------------------------------------|--|
| Standalone 2<br>(Patagonia) | Async <input type="radio"/> Sync <input type="radio"/> | Disable <input checked="" type="checkbox"/> Enable <input type="checkbox"/> |  |
| Standalone 2<br>(Darwin)    | Async <input type="radio"/> Sync <input type="radio"/> | Disable <input checked="" type="checkbox"/> Enable <input type="checkbox"/> |  |

1. Configure each standalone copy for consistency groups (or all standalone copies for group sets).

Standalone copies are not linked to the production, and you must decide how to handle them before failover. By default, RecoverPoint for VMs does not delete copy VMs but it does disable them. You can **Enable** any required standalone copies and select a replication mode (sync or async), or **Delete** them from the consistency group. Deleting a copy does not delete the VMs from storage.

**CAUTION: Disabled copy VMs require a full sweep when they are re-enabled.**

2. Click **SET AS PRODUCTION** to permanently failover.
  - o The role of the **Production** becomes **Remote/Local Copy**.
  - o The role of the **Remote/Local Copy** becomes **Production**.
  - o The standalone copies are handled as specified.
  - o The production VM and copy VM change roles, but their names do not change. Therefore, after failover, new production VMs will still be named *YourVMName\_copy* and the new copy VMs are still named *YourVMName*.
  - o The production journal becomes the copy journal and the copy journal becomes the production journal. The production journal does not contain the copy history, so it is by default, a much smaller journal. Therefore, after failover, when the production becomes the copy, you may want to add journal volumes to the new copy journal to ensure that you have ample space for copy testing. For detailed instructions on how to add journal volumes to a copy journal, see *Managing group protection policies*.
  - o The marking information in the production journal is deleted, the journal of the copy to which you failed over is deleted, and the consistency group undergoes a full sweep.

**CAUTION: During the full sweep, data is not transferred synchronously.**

## Manage production recovery

Manage the ongoing recovery of a copy of a consistency group or group set from the production using the **Recovery Activities** screen.

- NOTE:** When recovering of a copy from the production, wait for the **Activity Status** to show **Ready for next action** and the **Progress status bar** to reach 100%, indicating the specified snapshot image has been accessed.

When the **Activity Status** is **Ready for next action**:

- To select a consistency group for production recovery, ensure the **Consistency Groups** tab is selected.

#### Recovery Activities

- Click **ACTIONS > Start new test** to select another snapshot to test, or to redefine the testing network.
- Click **ACTIONS > Stop activity** to roll back all writes that were made to the copy snapshot while write access was enabled to the snapshot volumes, and disable write access to the snapshot volumes.
- (Optional) Click **UNDO WRITES** to undo all writes that were made to the copy snapshot while write access was enabled, without disabling access to the copy volumes.
- (Optional) Click **ENABLE DIRECT ACCESS** to write directly to the copy storage. Any changes made to the copy storage while directly accessing the copy cannot be automatically undone, because when a snapshot is directly accessed, the journal at the copy is deleted. On the other hand, direct access does not impose a limit to the amount of data that you can write to the copy storage volumes.

**CAUTION:** When direct access is enabled, replication stops to the copy, and a short initialization is required across all group volumes when direct access is disabled.

- NOTE:** The **Image Access Log Capacity** progress bar indicates how long you can access the copy image before the image access log is full and all writes to the copy fail. If you need more time for testing, you can undo writes at the copy, enable direct access to the copy, or add journal volumes as described in Managing group protection policies.

- To select a group set for production recovery, click the **Group Sets** tab.

#### Recovery Activities

- Click **OPEN** to display the **Detailed Status** of all consistency groups in the group set. After access is enabled to the copy snapshot, the **Status** column of all groups displays **Ready for next action**.

**NOTE:** Groups in the group set without a copy at the specified vRPA cluster are excluded from the activity.

## Detailed Status of 'group-set'



| Consistency Groups | Copy       | Snapshot            | Status                | Progress      |
|--------------------|------------|---------------------|-----------------------|---------------|
| cg_Win_206         | Local Copy | May 26, 2020 2:04:2 | Ready for next action | ████████ 100% |
| cg_Win             | Standalone | May 26, 2020 2:06:1 | Ready for next action | ████████ 100% |

Items per page: 10 — 2 Detailed Statuses

CLOSE

- Click **ACTIONS > Stop activity** to roll back all writes that were made to the copy snapshot while access was enabled to the copy volumes, and disable access to the copy volumes.

To recover production from the copy click **ACTIONS > Recover production**.

### NOTE:

- Data transfer from the production to all copies is paused, and will resume only after production recovery is complete.
- Host access to the recovered production volumes, and the recovering copy volumes is blocked.
- Recovered production volumes are overwritten. Any writes made to the copy during testing are transferred to the production, unless you clicked **UNDO WRITES**.
- The group undergoes a short initialization process to synchronize the new production data at the copy.

## Configuring email alerts and reports

Run a series of **Sysmgmt CLI** commands to configure your system to send system alerts and reports to a specified email, in real time.

### Prerequisites

Create an SSH connection to a vRPA cluster management IP address, and use your RecoverPoint for VMs **admin** username and password to log into the **Boxmgmt CLI > Sysmgmt CLI**.

### About this task

- NOTE:** For more information, refer to the *RecoverPoint for VMs CLI Reference Guide*.

### Steps

- Set up email access using the Sysmgmt CLI:
  - Run the **set\_smtp\_server** command, and enter the IP address or DNS name for sending email notifications. You should receive confirmation that the SMTP server has been configured successfully.
  - Run the **config\_email** command, and provide the requested information. You should receive confirmation that your email mechanism has been successfully configured.
  - Run the **enable\_email** command, and choose the **enable\_email** option. You should receive confirmation that email alerts have been enabled successfully.
- (Optional) To add email users, run the **add\_email\_users** command, and provide the requested information. You should receive confirmation that email users have been added successfully.
- (Optional) To configure system reports, in the Sysmgmt CLI, run the **config\_system\_reports** command, and provide the requested information. You should receive confirmation that system notifications have been successfully configured.

**Results**

Your email alerts and reports are configured.

**Next steps**

See the Monitor system alerts for more information.

# Troubleshooting

Use the following information, features and tools to troubleshoot your RecoverPoint for VMs.

## Topics:

- Finding the vRPA cluster management IP
- Collecting logs
- Adding new VMDKs
- Removing a VMDK
- Automatically expanding copy VMDKs
- Recovering from a cluster disaster
- RecoverPoint for VMs licensing
- Register RecoverPoint by email or phone
- Creating VMkernel ports
- Load balancing
- Copy VM network configuration guidelines
- Changing the network adapter configuration of a protected VM

## Finding the vRPA cluster management IP

Displays the vRPA cluster management IP of a specific vRPA cluster.

### Steps

1. In the vSphere HTML5 plugin, select **System > Administration**, and the **vRPA Clusters** tab.
2. Select the vRPA cluster.
3. Note the **Management IP Address** for the selected vRPA cluster.

## Collecting logs

Collecting logs is relevant only in support cases, and should be performed only when instructed to do so by Customer Support.

### Collecting logs from vRPA clusters

In RecoverPoint for VMs 5.3 SP2 and later versions, you can initiate the log collection process from within the RecoverPoint for VMs vSphere plugin.


#### About this task

Logs can be collected from one or more vRPAs in multiple vRPA clusters, as long as they all reside on a vCenter Server registered with the plugin server, or a vCenter Server linked to a vCenter Server that is registered with the plugin server.

**NOTE:** In RecoverPoint for VMs versions prior to 5.3 SP2, see the *RecoverPoint for VMs Installation and Deployment* Guide "Collect logs" section for information on how to collect logs from vRPA clusters.

### Steps

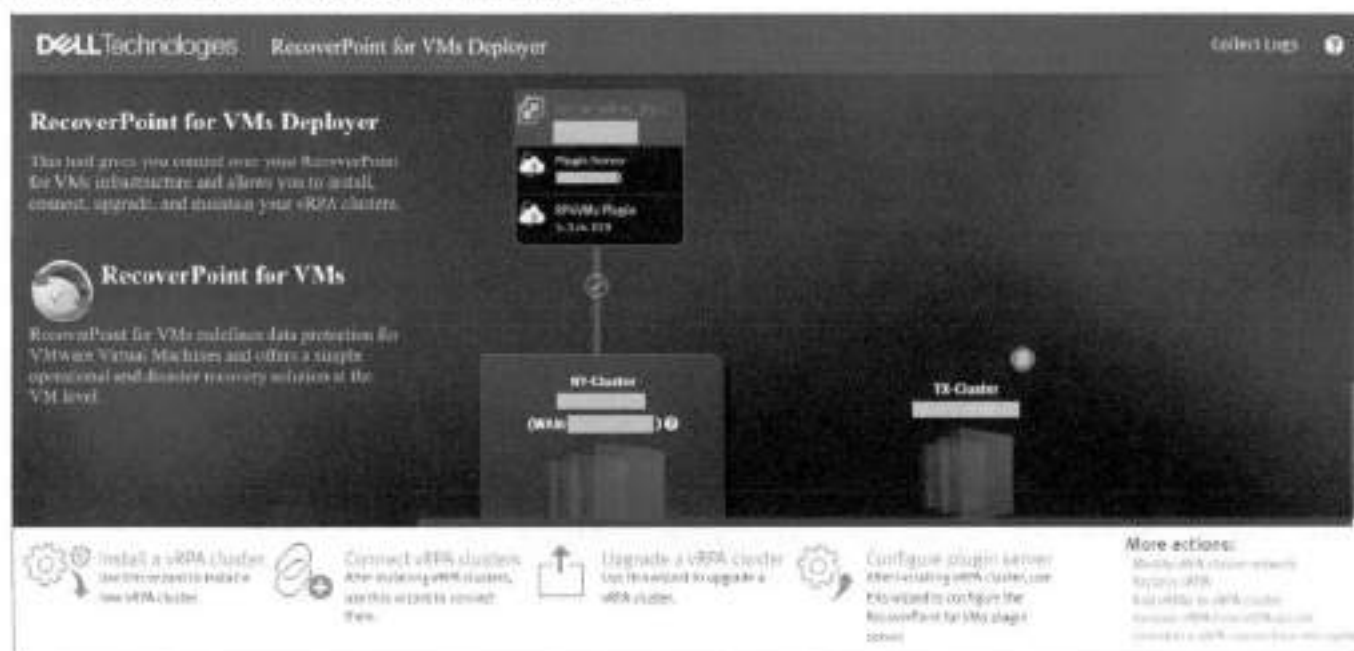
1. Click **System > Administration > vRPA Clusters**.
2. Select a vRPA cluster and click the **Collect logs** icon.



3. In the system notification that is displayed, click **Copy**, open a browser window and paste the copied URL into the browser address bar.



4. If prompted, type the login credentials for the **admin** user and click **Sign in**.
5. In the **RecoverPoint for VMs Deployer**, click **Collect Logs**.



The **Collect Cluster Logs** dialog is displayed.

## Collect Cluster Logs

Collect logs from vRPA cluster VxRail\_Site3

Enter time frame for collecting logs

Start time 07/08/2021 08:34 AM

End time 07/08/2021 12:34 PM

### Advanced

Collect logs from vRPAs in other vRPA clusters (collects full logs)

Collect Logs

Cancel

In the **Collect Cluster Logs** dialog:

- Enter a time frame for log collection.
- (Optional) Enter the IP addresses of other vRPA clusters from which to collect logs in the **Advanced** section.
- Click **Collect Logs**.

### Results

Depending on the size of the environment, log collection may take several minutes to complete. When the collection process is complete, a success message is displayed with the location (i.e. vRPA cluster) containing the logs.

Logs collected successfully for:

| Location | Log file                                                  |
|----------|-----------------------------------------------------------|
| Site1    | sysinfo-incomplete-Site1_KBox-1-2-2021.07.12.15.29.35.tar |

To download the logs, please use your 'admin' credentials

### Next steps

- In the success message, click the name of a vRPA cluster to open a browser window to the location of the collected logs.
- If prompted to, log into the vRPA cluster with your **admin** user credentials.
- Click a vRPA log name to download the vRPA log.

The name of each vRPA log has a \*.tar extension and it includes the <clustername><vrpaname> and <vrpaip> for easy identification. The log collection date is displayed under **Last Modified**.

## Directory Listing For [//]

| Filename                                                  | Size        | Last Modified                 |
|-----------------------------------------------------------|-------------|-------------------------------|
| is_report                                                 | 4.4 kb      | Mon, 12 Jul 2021 15:32:29 GMT |
| sysinfo-incomplete-site1_kbox-1-2-2021.07.12.15.29.35.tar | 198368.0 kb | Mon, 12 Jul 2021 15:32:29 GMT |
| long_term_stats/                                          |             | Mon, 12 Jul 2021 15:30:49 GMT |

## Collect plugin server logs

Collect plugin server logs for support purposes.

### About this task

The procedure for collecting plugin server logs is detailed in *Collecting plugin server logs*.

## Collecting RecoverPoint for VMs splitter logs

### About this task

RecoverPoint for VMs splitter logs are in the ESXi logs. To export the ESXi system logs, use the following procedure.

### Steps

1. In the vSphere Client, select **Menu > Home and Clusters**.
2. Right click on the desired vCenter, and select **Export System Logs...**
3. In the **Select hosts** pane of the **Export System Logs** screen:
  - a. Select the ESXi hosts for which you want to export the system logs.
  - b. (Optional) Select **Include vCenter and vSphere UI Client logs**.
  - c. Select the system logs to be exported
  - d. (Optional) Select **Gather performance data**, and specify a duration and interval
  - e. (Optional) Set a password with which to encrypt the collected log data.
4. In the **Select logs** pane of the **Export System Logs** screen:
  - a. Select the system logs to be exported
  - b. (Optional) Select **Gather performance data**, and specify a duration and interval
  - c. (Optional) Set a password with which to encrypt the collected log data.
5. Click **EXPORT LOGS**.

## Adding new VMDKs

Best practices and system behavior when adding new VMDKs to a protected VM.

### About this task

RecoverPoint for VMs automatically detects when a new VMDK is added to a protected VM through the vSphere Client VM Properties, and by default, automatically starts protecting each added VMDK.

- To disable automatic protection for all VMDKs added to a protected VM in the future, see *Automatic protection of newly added VMDKs*.
- To exclude specific VMDKs of a protected VM from protection, see *Excluding a VMDK from replication*.

### Results

When a new VMDK is added to a protected VM, a volume sweep occurs on the added VMDK and a short initialization occurs on all other VMDKs in the consistency group, but no history is lost.

## Removing a VMDK

Defines how to handle a VMDK which was removed from a protected VM, at the copy.

### About this task

RecoverPoint for VMs automatically detects when a VMDK is removed from a protected VM through the vSphere Client, and displays an alert when there is hardware mismatch between a protected VM and its copy VM.



- If the production VMDK removal was intentional, follow the instructions in [Excluding a VMDK from replication](#) to stop protecting it.
- If the production VMDK removal was unintentional, run [Recover production from a copy](#) to recover the removed VMDK. For recover production, select a snapshot that pre-dates VMDK removal.

**NOTE:** Removing VMDKs from a protected VM does not delete their copies and does not remove their history from the copy journal.

## Automatically expanding copy VMDKs

Best practices, troubleshooting, system behavior, and limitations of automatically expanding copy VMs when a protected VM is expanded.

### About this task

When a protected VMDK is expanded, RecoverPoint for VMs automatically expands all corresponding copy VMDKs, with the following limitations:

- VMDKs can be expanded, but they cannot be shrunk.
- When a production VMDK is expanded, the system pauses replication of the consistency group while the system is busy resizing the corresponding copy VMDK.
- Automatic VMDK expansion fails if:
  - Replicating to RDM. After expanding the production VMDK or RDM, you must manually expand the copy RDM.
  - The datastore does not contain enough free space, and you should free up space in the copy VM datastore.
  - A snapshot has been taken of the virtual machine containing the copy VMDK. Enable access to the copy containing the VMDK and use the vCenter snapshot manager to delete all snapshots before disabling image access.
  - The version of the file system that you are running does not support the VMDK size. In this case, consider upgrading the file system version.
  - The size of a copy VMDK is larger than the size of its corresponding production VMDK. In this case, to begin the automatic VMDK expansion process, you must manually expand the production VMDK. This manual expansion might be required if you failed over while automatic expansion was in progress, or if the copy VMDK was manually expanded.
  - Replicating to RDM. After expanding the production VMDK or RDM, you must manually expand the copy RDM for replication to resume.
  - A snapshot of a copy containing a VMDK marked for automatic expansion is selected during testing or recovery. In this case, you should disable image access for replication to resume.
  - A protected VMDK is smaller than the size registered in the system settings. In this case you should contact Customer Support. This can happen, for example, if a production VMDK has been removed and re-added with a smaller size.
  - One or more copy VMDKs have been marked for automatic expansion, but the system cannot automatically resize a RAW device. In this case, enable access to the copy VM with the problematic VMDK and manually expand it before disabling image access. If problem persists, contact Customer Support.

### Results

After fixing any of these issues, wait 15 minutes for the automatic expansion process to restart and the error to resolve itself. If the problem persists, try manually resizing the copy VMDKs or contact Customer Support.

## Recovering from a cluster disaster

After a full cluster disaster or a switch disaster, it may take 10 minutes or more for all the components of the vRPA system to restart, reconnect, and restore full operation.

# RecoverPoint for VMs licensing

RecoverPoint for VMs supports two types of licensing models: VM-based licensing and socket-based licensing.

## VM-based licensing

With VM-based licensing, licenses are based on the number of supported VMs per vCenter Server. Only production VMs are counted in the number of supported VMs per vCenter Server. Licensing is enforced using the vCenter Server ID.

All vCenter Servers must be registered in RecoverPoint for VMs before their licenses can be added. vCenter Server registration is performed in the RecoverPoint for VMs Deployer UI. Refer to the *RecoverPoint for VMs Installation and Deployment Guide* for more information.

You can use the HTML5 plugin to register additional vCenter Servers. For details, see *Register vCenter Server to vRPA cluster*.

When you reach the maximum number of VMs that the license supports for each vCenter Server, you cannot protect new VMs or enable disabled consistency groups. However, replication of existing VMs and consistency groups continues.

Failover has no effect on the license.

## Socket-based licensing

With socket-based licensing, licenses are based on the number of physical CPU sockets in the ESXi hosts that host the production VMs. A VM does not 'belong' to a specific socket.

When you reach the maximum number of sockets that the license supports for each vCenter Server, you cannot protect new VMs or enable disabled consistency groups. However, replication of existing VMs and consistency groups continues.

As with VM-based licensing, failover does not affect the socket-based license. However, vMotion of production VMs does affect the license and may cause a license violation due to an increase in the number of sockets being used. ESXi hosts that host the production VMs are the ones that count in a socket-based license. To avoid license violations, it is a best practice to license all ESXi hosts of the ESXi cluster.

## Adding a socket-based license to a system with VM-based licenses

When using VM-based licensing, license capacity is measured by the number of VMs. For example, when you view the license capacity in the UI, it may be listed as:

```
Capacity = 30 VMs
```

When using socket-based licensing, license capacity is measured by the number of sockets. For example, the license capacity may be listed as:

```
Capacity = 2 sockets
```

When a socket-based license is installed on a RecoverPoint for VMs system that has VM-based licenses, the system automatically converts VM-based licenses to socket-based licenses at a ratio of 15 VMs per socket. In this case, the license capacity would be listed as:

```
Capacity = 30 VMs (2 sockets)
```

In cases where the ratio does not result in an even conversion, the value is rounded up. For example:

```
Capacity = 31 VMs (3 sockets)
```

Since licenses are applied per vCenter, and not per vRPA cluster, multiple vRPA clusters with VMs or CPU sockets may count towards the same license.

## License subscriptions

VM- and socket-based licenses may be installed as subscriptions. Unlike a permanent license, a subscription license has a start date and an end date. The system sends an alert beginning 30 days before license expiration to indicate the number of days remaining. Subscription and permanent licenses may coexist.

You can install a subscription license before its start date. It automatically becomes active on the start date.

## Register RecoverPoint by email or phone

If your company is without external connectivity, and you cannot register your RecoverPoint for VMs system online, you can also register by phone.

### About this task

- Register the RecoverPoint system after:
  - Installing a RecoverPoint system
  - Connecting RPA clusters in a RecoverPoint system
  - Upgrading a RecoverPoint system
- The registration process is incomplete if valid values are not provided for every field in the post-deployment form.

### Steps

1. Gather the required information.
  - Download the post-deployment form:
    - a. Access <https://www.dell.com/support>
    - b. Search for the term *Post-Deployment Form*
  - If you have access to a Flex plugin, fill out the RecoverPoint and RecoverPoint for VMs Post-Deployment Form, for every vRPA cluster
  - Export the RecoverPoint registration information, for every vRPA cluster. Using the Flex plugin:
    - a. Select **Administration > vRPA Clusters**.
    - b. Select the vRPA cluster for which you want to export a post-deployment form, and then click **Support**.
    - c. In the Registration pane, click the **Export to CSV** button and save the file to the computer.
2. Send the information to the Install Base group:
  - Customers and partners: Email the post-deployment form to the Install Base group at [rp.registration@emc.com](mailto:rp.registration@emc.com).
  - Employees:
    - (Preferred) Use the Install Base Group under Post Sales at <http://emc.force.com/BusinessServices>.
    - Call in the information to the Install Base Group at 1-866-436-2411 – Monday to Friday (normal Eastern Time Zone working hours).

## Creating VMkernel ports

If before clicking protecting VMs you received a warning regarding a potential communications problem, and after clicking **Protect**, transfer for the consistency group does not eventually reach the **Active** status, you may need to create VMkernel ports for all ESXi hosts in the cluster.

### Steps

1. Select **System > ESX Clusters**, and click the **Configure VMkernel ports** icon of an ESXi cluster.



- In the **VMkernel Port Settings** dialog box, specify the settings, including a range of available IPs, for creating VMkernel ports for all ESXi hosts in the cluster.

## VMkernel Port Settings

**VIRTUAL SWITCH** DISTRIBUTED VIRTUAL SWITCH

### Virtual Switch

vSwitch1

### IP Pool Range

192.168.0.21#4

Enter IP address ranges as an ordered, comma-separated list, for example: 1.2.3.4#70, 1.2.3.80#16

### Subnet Mask

255.255.255.0

Most likely same as the vRPA data netmask

### VLAN

0

### MTU

1500

CANCEL

**DONE**

- Click **DONE**

## Load balancing

### About this task

Load balancing is the process of assigning preferred vRPAs to consistency groups so that the preferred vRPA performs data transfer for that group. This is done to balance the load across the system and to prevent the system from entering a high-load state.

Perform load balancing:

- When a new consistency group is added to the system. Wait 1 week after the new group is added to accumulate enough traffic history before performing load balancing.
- When a new vRPA is added to a vRPA cluster. Perform load balancing immediately after the vRPA is added.

- If the system enters high load frequently. When load balancing is required, the event logs display a message indicating so. When you see this message, perform load balancing.
- Periodically, to ensure that the system is always handling distributing loads evenly. A script can be created to periodically perform load balancing.

### Steps

1. To balance the load on the vRPAs, use an `ssh` client to connect to the vRPA management IP address, and type the `RecoverPoint` username and password to log in to the CLI.
2. Run the `balance_load` command to balance the load. To view command parameters that can refine the search, run: `balance_load ?`

## Copy VM network configuration guidelines

Use the following guidelines for *Re-IP* rules.

**Table 1. Virtual machine network settings available through the GUI**

| Setting                      | Description                                                                                                         | Guidelines                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------|---------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>(VM) Operating System</b> | The guest operating system of the specified VM.                                                                     | <ul style="list-style-type: none"> <li>• Not customizable.</li> <li>• Automatically populated by the system.</li> <li>• Possible values are <i>Windows</i>, <i>Linux</i>, or <i>Unknown</i>.</li> </ul>                                                                                                                                                                                                                                                                                                      |
| <b>(VM) Host Name</b>        | The hostname of the specified VM.                                                                                   | <ul style="list-style-type: none"> <li>• Only mandatory for virtual machines with a Linux operating system.</li> <li>• Customizable.</li> <li>• Value can be retrieved from production VM using <b>Get Value from Production</b> or <b>Get All Values from Production</b>.</li> </ul>                                                                                                                                                                                                                        |
| <b>(VM) DNS Domain</b>       | The DNS domain for the specified VM.                                                                                | <ul style="list-style-type: none"> <li>• Only relevant (and mandatory) for virtual machines with a Linux operating system.</li> <li>• Value should be in the format <b>example.company.com</b>.</li> </ul>                                                                                                                                                                                                                                                                                                   |
| <b>(VM) DNS Server(s)</b>    | The global IP address that identifies one or more DNS servers for all adapters of the specified VM.                 | <ul style="list-style-type: none"> <li>• Only relevant for virtual machines with a Linux operating system.</li> <li>• Customizable.</li> <li>• Can be left blank.</li> <li>• This setting applies to all virtual network adapters of the specified VM.</li> <li>• Separate multiple values with a semicolon (;).</li> <li>• Value can be retrieved from production VM using <b>Get Value from Production</b> or <b>Get All Values from Production</b>.</li> </ul>                                            |
| <b>(VM) DNS Suffix(s)</b>    | The global settings of the suffixes for the DNS servers of all adapters on both Windows and Linux virtual machines. | <ul style="list-style-type: none"> <li>• Customizable.</li> <li>• Can be left blank.</li> <li>• Value can be retrieved from production VM using <b>Get Value from Production</b> or <b>Get All Values from Production</b>.</li> </ul>                                                                                                                                                                                                                                                                        |
| <b>(Adapter) IP Address</b>  | IPv4 address for this virtual network adapter.                                                                      | <ul style="list-style-type: none"> <li>• Can contain either a static IPv4 address or DHCP string.</li> <li>• Can be left blank when using IPv6.</li> <li>• Define one IPv4 address, one IPv6 address, or one of each, for the same virtual network adapter. Entering multiple IPv4 or IPv6 addresses for the same virtual network adapter is not supported.</li> <li>• Value can be retrieved from production VM using <b>Get Value from Production</b> or <b>Get All Values from Production</b>.</li> </ul> |

Table 1. Virtual machine network settings available through the GUI (continued)

| Setting                                    | Description                                                               | Guidelines                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------------------------|---------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>(Adapter) Subnet</b>                    | IPv4 subnet mask for this virtual network adapter.                        | <ul style="list-style-type: none"> <li>• Mandatory when an <b>IP Address</b> is entered.</li> <li>• Can be left blank when using IPv6.</li> <li>• Value can be retrieved from production VM using <b>Get Value from Production</b> or <b>Get All Values from Production</b>.</li> </ul>                                                                                                                                                                                                                           |
| <b>(Adapter) Gateway(s)</b>                | One or more IPv4 gateways for this virtual network adapter.               | <ul style="list-style-type: none"> <li>• Mandatory when an <b>IP Address</b> is entered.</li> <li>• Can be left blank when using IPv6.</li> <li>• Separate multiple values with a semicolon (;).</li> <li>• Value can be retrieved from production VM using <b>Get Value from Production</b> or <b>Get All Values from Production</b>.</li> </ul>                                                                                                                                                                 |
| <b>(Adapter) IPv6 Address</b>              | IPv6 address for this virtual network adapter.                            | <ul style="list-style-type: none"> <li>• Can contain either a static IPv6 address or it's DHCP string.</li> <li>• Can be left blank when using IPv4.</li> <li>• Define one IPv4 address, one IPv6 address, or one of each, for the same virtual network adapter. Entering multiple IPv4 or IPv6 addresses for the same virtual network adapter is not supported.</li> <li>• Value can be retrieved from production VM using <b>Get Value from Production</b> or <b>Get All Values from Production</b>.</li> </ul> |
| <b>(Adapter) IPv6 Subnet Prefix Length</b> | IPv6 subnet mask for this virtual network adapter.                        | <ul style="list-style-type: none"> <li>• Customizable.</li> <li>• Can be left blank when using IPv4.</li> <li>• Value can be retrieved from production VM using <b>Get Value from Production</b> or <b>Get All Values from Production</b>.</li> </ul>                                                                                                                                                                                                                                                             |
| <b>(Adapter) IPv6 Gateway(s)</b>           | One or more IPv6 gateways for this virtual network adapter.               | <ul style="list-style-type: none"> <li>• Customizable.</li> <li>• Mandatory when an IPv6 format <b>IP Address</b> is entered.</li> <li>• Can be left blank when using IPv4.</li> <li>• Separate multiple values with a semicolon (;).</li> <li>• Value can be retrieved from production VM using <b>Get Value from Production</b> or <b>Get All Values from Production</b>.</li> </ul>                                                                                                                            |
| <b>(Adapter) DNS Server(s)</b>             | IP address of one or more DNS server(s) for this virtual network adapter. | <ul style="list-style-type: none"> <li>• Can be left blank.</li> <li>• Can contain one or more IPv4 DNS servers for each virtual network adapter (NIC).</li> <li>• Applies only to the configured adapter when a value other than <b>Adapter ID 0</b> is defined.</li> <li>• Separate multiple values with a semicolon (;).</li> <li>• Value can be retrieved from production VM using <b>Get Value from Production</b> or <b>Get All Values from Production</b>.</li> </ul>                                      |
| <b>(Adapter) NetBIOS</b>                   | Whether or not to activate NetBIOS on this virtual network adapter.       | <ul style="list-style-type: none"> <li>• Cannot be left blank.</li> <li>• Only relevant for virtual machines running a Windows operating system.</li> <li>• Default is <b>Enabled</b>.</li> <li>• Net-BIOS should be enabled.</li> <li>• Valid values are <b>DISABLED</b>, <b>ENABLED</b>, <b>ENABLED_VIA_DHCP</b>.</li> </ul>                                                                                                                                                                                    |
| <b>(Adapter) Primary WINS</b>              | Primary WINS server of this virtual network adapter.                      | <ul style="list-style-type: none"> <li>• Relevant for windows virtual machines only.</li> <li>• Customizable.</li> <li>• Can be left blank.</li> </ul>                                                                                                                                                                                                                                                                                                                                                            |

Table 1. Virtual machine network settings available through the GUI (continued)

| Setting                  | Description                                            | Guidelines                                                                                                                                       |
|--------------------------|--------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| (Adapter) Secondary WINS | Secondary WINS server of this virtual network adapter. | <ul style="list-style-type: none"> <li>Relevant for windows virtual machines only.</li> <li>Customizable.</li> <li>Can be left blank.</li> </ul> |

Table 2. Network settings only available through the JSON file

| Setting              | Description                                                                                                                         | Guidelines                                                                                                                                                                                                                                                                               |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CG ID                | The consistency group ID in the RecoverPoint for VMs system.                                                                        | <ul style="list-style-type: none"> <li>Do not modify this field.</li> <li>Automatically populated by the system.</li> <li>Not customizable.</li> <li>Can be left blank.</li> </ul>                                                                                                       |
| CG Name              | Name of the consistency group in the RecoverPoint for VMs system.                                                                   | <ul style="list-style-type: none"> <li>Automatically populated by the system.</li> <li>Must be the name associated with the specified consistency ID in RecoverPoint for VMs.</li> <li>Customizable.</li> <li>Can be left blank.</li> </ul>                                              |
| VC ID                | The vCenter Server ID in VMware.                                                                                                    | <ul style="list-style-type: none"> <li>Do not modify this field.</li> <li>Automatically populated by the system.</li> <li>Not customizable.</li> <li>Can be left blank.</li> </ul>                                                                                                       |
| VC Name              | The name of the vCenter Server hosting the virtual machine.                                                                         | <ul style="list-style-type: none"> <li>Customizable.</li> <li>Can be left blank.</li> </ul>                                                                                                                                                                                              |
| VM ID                | The virtual machine ID that vCenter Server uses.                                                                                    | <ul style="list-style-type: none"> <li>Do not modify this field.</li> <li>Automatically populated by the system.</li> <li>Not customizable.</li> <li>Cannot be left blank.</li> </ul>                                                                                                    |
| VM Name              | The name of the virtual machine.                                                                                                    | <ul style="list-style-type: none"> <li>Customizable.</li> <li>Automatically populated by the system.</li> <li>Can be left blank.</li> </ul>                                                                                                                                              |
| NIC Index in vCenter | The index of the adapter in the order of virtual network adapters (NICs) in the virtual machine settings of the vCenter web client. | <ul style="list-style-type: none"> <li>Customizable.</li> <li>Cannot be left blank.</li> <li>Enter a numeric value.</li> <li>Enter a value of 0 to define the first virtual network adapter in the vSphere Web Client. Enter a value of 1 to define the next network adapter.</li> </ul> |

## Changing the network adapter configuration of a protected VM

### About this task

When the virtual network adapter (NIC) configuration of a production VM changes, any pre-existing copy VM network configuration may be adversely affected and may require re-configuration before it works. After adding or removing NICs from a protected virtual machine, re-configure the copy VM network using Re-IP rules.

If the NIC configuration of a production VM changes and the change is not reflected in the copy VM, ensure **Hardware changes** is enabled and *enable image access by Test a copy*.

## Shared disks

Shared disks are VMDKs or RDMs that are mapped to and accessible from multiple VMs. They are commonly used with database systems such as Oracle RAC and Microsoft SQL Server, where shared disks hold the database and all nodes in the cluster access it.

Use the procedures in this appendix to manage systems that use shared VMDKs or RDMs.

### Topics:


- Add VM to a host cluster
- Add shared disk to VMs that belong to a consistency group
- Recover production after deletion of production VM or shared disk
- Failover after deletion of production VM or shared disk
- Remove a VM from its host cluster
- Unprotect a shared disk

## Add VM to a host cluster

Use this procedure to add a VM to a Microsoft Failover Cluster or Oracle RAC whose VMs are already protected by RecoverPoint for VMs.

### Steps

1. Add all shared disks to the new production VM.
2. Create a VM with the same configuration as the copy.
3. Add all shared disks at the copy to the new copy VM.
4. Using the Flex plugin, add the new VM to the existing consistency group. You must manually select the copy VM that you created in Step 2.

 **NOTE:** You can manually select the copy VM only when using the Flex plugin.

5. Add the newly protected VM to the Microsoft Failover Cluster or Oracle RAC.

### Results

The newly added VM is now protected.

## Add shared disk to VMs that belong to a consistency group

Use this procedure to add a shared disk to all VMs that belong to an existing consistency group.

### Steps

1. Add a shared disk (VMDK or RDM) to the production VMs that belong to the host Microsoft Failover Cluster or Oracle RAC.
2. Add a shared disk with the same configuration to the parallel VMs at the copy.
3. Use the RecoverPoint for VMs plugin to protect the new shared disk by including the shared disk on all of the production VMs.
4. Add the shared disk to the Microsoft Failover Cluster or Oracle RAC.



### Results

The new shared disk is added to all of the VMs that belong to the consistency group.

## Recover production after deletion of production VM or shared disk

Use this procedure to restore a shared disk (VMDK or RDM) that has been removed from production VMs.

### About this task

RecoverPoint for VMs automatically recreates removed VMs, but does not recreate removed shared disks.

### Steps

1. If the shared disk is removed from all of the VMs that were mapped to it, then recreate the shared disk. If one or more of the VMs remain, then the shared disk is not deleted. Do not create a new shared disk.
2. Recover production to a selected image.  
Recover production will not proceed if you have not recreated the removed shared disk.
3. After production recovery, to resume replication from production to copy, remap the shared disk to all of the VMs to which it was previously mapped.

### Results

The shared disk is restored, and replication can resume.

## Failover after deletion of production VM or shared disk

Use this procedure for failover of a consistency group to a copy following removal of a production VM or shared disk (VMDK or RDM).

### About this task

RecoverPoint for VMs automatically recreates removed VMs, but does not recreate removed shared disks.

### Steps

1. Failover the consistency group to a copy.  
Even when production VMs or shared disks are removed, it is still possible to successfully failover the consistency group.
2. After failover, to resume replication from the copy to production, restore the removed VMs and shared disks at production.  
If one or more of the VMs remain, then the shared disk is not deleted. Do not create a new shared disk. Ensure that the shared disk is added to all of the VMs that were previously mapped to it.  
If the shared disk at production is removed from all of the VMs that were mapped to it, then recreate the shared disk, and add it to all of the VMs that were previously mapped to it.

### Results

Failover to the copy is successful, and data replication can resume from the copy to production.

## Remove a VM from its host cluster

Use this procedure to unprotect a VM that belongs to a Microsoft Failover Cluster or Oracle RAC.

### Steps

1. Remove a VM from its host cluster.
2. Unmap all shared disks from the VM.
3. Remove the VM from its consistency group.

### Results

The VM is removed from its host cluster.

## Unprotect a shared disk

Use this procedure to unprotect a shared disk (VMDK or RDM).

### Steps

1. Use the RecoverPoint for VMs plugin to exclude from RecoverPoint for VMs protection all VMs that are mapped to the shared disk .
2. For any shared disk that is removed from the production VMs, remove the parallel shared disks from the copy VMs. Snapshots taken during the time when the shared volume is partially protected are inconsistent.

### Results

The shared disk is removed from protection.



# Dell vProtect 19.14.0.600 Compatibility

|                                                                                    |           |
|------------------------------------------------------------------------------------|-----------|
| <b>DEPLOYMENT</b> .....                                                            | <b>3</b>  |
| <i>Deployment (Physical or Virtual Machine)</i> .....                              | 3         |
| <i>Deployment (VMware Virtual Appliance)</i> .....                                 | 3         |
| <i>Deployment (Red Hat Virtualization   oVirt   OLVM Virtual Appliance)</i> .....  | 3         |
| <i>Deployment (Citrix Hypervisor – XenServer   XCP-ng Virtual Appliance)</i> ..... | 4         |
| <i>Deployment (Acropolis Hypervisor AHV Virtual Appliance)</i> .....               | 5         |
| <b>BACKUP TARGETS</b> .....                                                        | <b>5</b>  |
| <b>SUPPORTED PLATFORMS   VIRTUAL ENVIRONMENTS</b> .....                            | <b>6</b>  |
| <i>Nutanix Acropolis Hypervisor (AHV)</i> .....                                    | 6         |
| <i>Red Hat Virtualization   RHV</i> .....                                          | 7         |
| <i>oVirt</i> .....                                                                 | 8         |
| <i>Oracle Linux Virtualization Manager   OLVM</i> .....                            | 9         |
| <i>Proxmox VE</i> .....                                                            | 10        |
| <i>KVM/Xen Legacy</i> .....                                                        | 11        |
| <i>OpenStack Community   Vanilla</i> .....                                         | 12        |
| <i>Red Hat OpenStack</i> .....                                                     | 13        |
| <i>RDO OpenStack</i> .....                                                         | 14        |
| <i>Ubuntu OpenStack</i> .....                                                      | 15        |
| <i>Virtuozzo</i> .....                                                             | 16        |
| <i>Oracle VM</i> .....                                                             | 17        |
| <i>Citrix Hypervisor (formerly Xen Server)</i> .....                               | 18        |
| <i>XCP-ng</i> .....                                                                | 19        |
| <i>Microsoft Hyper-V</i> .....                                                     | 20        |
| <i>Azure Stack HCI</i> .....                                                       | 21        |
| <i>Huawei FusionCompute</i> .....                                                  | 22        |
| <i>Scale Computing HC3</i> .....                                                   | 23        |
| <b>SUPPORTED PLATFORMS   STORAGE PROVIDERS</b> .....                               | <b>24</b> |

|                                    |    |
|------------------------------------|----|
| <i>Filesystem</i> .....            | 24 |
| <i>Ceph RBD</i> .....              | 26 |
| <i>Nutanix Files</i> .....         | 27 |
| <i>Nutanix Volume Groups</i> ..... | 28 |

## Deployment

### Deployment (Physical or Virtual Machine)

| DEPLOYMENT (physical or virtual machine) |                                             |                                     |
|------------------------------------------|---------------------------------------------|-------------------------------------|
| Platform                                 | Platform Version                            | Comments                            |
| RHEL                                     | 8.0, 8.1, 8.2, 8.3, 8.4, 8.5, 8.6, 9.0, 9.1 | For general RPM based installation. |
| CentOS                                   | Stream 8, Stream 9                          | For general RPM based installation. |

### Deployment (VMware Virtual Appliance)

| DEPLOYMENT (VMware)                                                               |                  |                                                                                              |
|-----------------------------------------------------------------------------------|------------------|----------------------------------------------------------------------------------------------|
| Platform                                                                          | Platform Version | Comments                                                                                     |
| VMware vCenter/ESXi                                                               | 6.5, 6.7, 7.0    | Version 6.5 and later.                                                                       |
| VMware Tools                                                                      | 10 and later     | Version 10.1 and later is an optional to provide better experience with VM.                  |
| Web browser                                                                       | Google Chrome    | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |
| Notes:                                                                            |                  |                                                                                              |
| <ul style="list-style-type: none"><li>Supported datastores: vVols, vSAN</li></ul> |                  |                                                                                              |

### Deployment (Red Hat Virtualization | oVirt | OLVM Virtual Appliance)

| DEPLOYMENT (Red Hat Virtualization)    |                    |                                                                                              |
|----------------------------------------|--------------------|----------------------------------------------------------------------------------------------|
| Platform                               | Platform Version   | Comments                                                                                     |
| Red Hat Virtualization Manager (RHV-M) | 4.1, 4.2, 4.3, 4.4 | Version 4.1 and later.                                                                       |
| Guest Tools (oVirt guest agent)        | Current            | Required only for VM FS freeze                                                               |
| Web browser                            | Google Chrome      | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |

| DEPLOYMENT (oVirt)              |                         |                                                                                              |
|---------------------------------|-------------------------|----------------------------------------------------------------------------------------------|
| Platform                        | Platform Version        | Comments                                                                                     |
| oVirt                           | 4.1, 4.2, 4.3, 4.4, 4.5 | Version 4.1 and later.                                                                       |
| Guest Tools (oVirt guest agent) | Current                 | Required only for VM FS freeze                                                               |
| Web browser                     | Google Chrome           | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |

| DEPLOYMENT (Oracle Linux Virtualization Manager) |                  |                                                                                              |
|--------------------------------------------------|------------------|----------------------------------------------------------------------------------------------|
| Platform                                         | Platform Version | Comments                                                                                     |
| oVirt                                            | 4.2.8, 4.3, 4.4  | Version 4.2.8 and later.                                                                     |
| Guest Tools (oVirt guest agent)                  | Current          | Required only for VM FS freeze                                                               |
| Web browser                                      | Google Chrome    | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |

#### Deployment (Citrix Hypervisor – XenServer | XCP-ng Virtual Appliance)

| DEPLOYMENT (Citrix Hypervisor   XenServer) |                                   |                                                                                              |
|--------------------------------------------|-----------------------------------|----------------------------------------------------------------------------------------------|
| Platform                                   | Platform Version                  | Comments                                                                                     |
| Citrix Hypervisor (Xen Server) / XCP-ng    | 7.3, 7.4, 7.5, 7.6, 8.0, 8.1, 8.2 |                                                                                              |
| Citrix VM Tools                            | Current                           | Required only for VM FS freeze                                                               |
| Web browser                                | Google Chrome                     | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |

| DEPLOYMENT (XCP-ng)                     |                              |                                                                                              |
|-----------------------------------------|------------------------------|----------------------------------------------------------------------------------------------|
| Platform                                | Platform Version             | Comments                                                                                     |
| Citrix Hypervisor (Xen Server) / XCP-ng | 7.4, 7.5, 7.6, 8.0, 8.1, 8.2 |                                                                                              |
| Citrix VM Tools                         | Current                      | Required only for VM FS freeze                                                               |
| Web browser                             | Google Chrome                | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |

### Deployment (Acropolis Hypervisor AHV Virtual Appliance)

| DEPLOYMENT (Nutanix AHV) |                                                 |                                                                                              |
|--------------------------|-------------------------------------------------|----------------------------------------------------------------------------------------------|
| Platform                 | Platform Version                                | Comments                                                                                     |
| Acropolis Hypervisor AHV | 5.11, 5.15 LTS, 5.17, 5.18, 5.19, 5.20 LTS, 6.0 |                                                                                              |
| Nutanix Guest Tools      | Current                                         | Required only for VM FS freeze                                                               |
| Web browser              | Google Chrome                                   | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |

### Backup Targets

| BACKUP TARGETS                                                                                                                                   |                                                                         |                                                 |
|--------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------|-------------------------------------------------|
| Data Domain (Physical & Virtual)                                                                                                                 |                                                                         |                                                 |
| Platform                                                                                                                                         | Platform Version                                                        | Comments                                        |
| DDOS                                                                                                                                             | 7.1.0, 7.2.0, 7.3.0, 7.4.0, 7.5.0, 7.6.0, 7.7.0, 7.8.0, 7.9.0, 7.10.0   | Version 7.1.0 and later.                        |
| DDMC                                                                                                                                             | 7.1.0, 7.2.0, 7.3.0, 7.4.0, 7.5.0, 7.6.0, 7.7.0, 7.8.0, 7.9.0, 7.10.0   | Version 7.1.0 and later.                        |
| DDVE                                                                                                                                             | 7.1.0, 7.2.0, 7.3.0, 7.4.0, 7.5.0, 7.6.0, 7.7.0, 7.8.0, 7.9.0, 7.10.0   | Version 7.1.0 and later.                        |
| Data Domain Boost – File System (BoostFS)                                                                                                        | 7.0.0.0, 7.2.0.5, 7.3.0.5, 7.4.0.5, 7.5.0.5, 7.6, 7.7, 7.8, 7.9, 7.10.0 | RHEL 8.0, 8.1, 8.2, 8.3, 8.4<br>CentOS Stream 8 |
| Notes:                                                                                                                                           |                                                                         |                                                 |
| <ul style="list-style-type: none"> <li>Dell EMC vProtect recommendation is to use min. package - DDBoostFS-7.0.0.0-633922.rhel.x86_64</li> </ul> |                                                                         |                                                 |



## Supported Platforms | Virtual Environments

### Nutanix Acropolis Hypervisor (AHV)

| Nutanix Acropolis AHV    |                                                                       |                                                                                              |
|--------------------------|-----------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| Platform                 | Platform Version                                                      | Comments                                                                                     |
| Acropolis Hypervisor AHV | 5.5, 5.9, 5.10, 5.11, 5.15 LTS, 5.16, 5.17, 5.18, 5.19, 5.20 LTS, 6.0 |                                                                                              |
| Nutanix Guest Tools      | Current                                                               | Required only for VM FS freeze                                                               |
| Web browser              | Google Chrome                                                         | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |

| VM File Level Recovery (Mountable Backups)           |           |                        |
|------------------------------------------------------|-----------|------------------------|
| Operating System                                     | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                        | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                        | Ubuntu    | XFS, EXT3, EXT4, BTRFS |

## Red Hat Virtualization | RHV

| Red Hat Virtualization   RHV           |                                     |                                                                                                                                     |
|----------------------------------------|-------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| Platform                               | Platform Version                    | Comments                                                                                                                            |
| Red Hat Virtualization Manager (RHV-M) | 3.5.1, 3.6, 4.0, 4.1, 4.2, 4.3, 4.4 | Export Storage Domain strategy - 3.x only<br>Disk image transfer/ SSH transfer strategy - 4.2 +<br>Disk attachment strategy - 4.0 + |
| Guest Tools (oVirt guest agent)        | Current                             | Required only for VM FS freeze                                                                                                      |
| Web browser                            | Google Chrome                       | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI.                                        |

| Red Hat Virtualization UI Plugin       |                  |          |
|----------------------------------------|------------------|----------|
| Platform                               | Platform Version | Comments |
| Red Hat Virtualization Manager (RHV-M) | 4.3, 4.4         |          |

| VM File Level Recovery (Mountable Backups)           |           |                        |
|------------------------------------------------------|-----------|------------------------|
| Operating System                                     | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                        | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                        | Ubuntu    | XFS, EXT3, EXT4, BTRFS |

## oVirt

| oVirt                           |                                          |                                                                                                                                         |
|---------------------------------|------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| Platform                        | Platform Version                         | Comments                                                                                                                                |
| oVirt                           | 3.5.1, 3.6, 4.0, 4.1, 4.2, 4.3, 4.4, 4.5 | Export Storage Domain strategy - 3.x only<br>Disk image transfer/ SSH transfer strategy/CBT - 4.3 +<br>Disk attachment strategy - 4.0 + |
| Guest Tools (oVirt guest agent) | Current                                  | Required only for VM FS freeze                                                                                                          |
| Web browser                     | Google Chrome                            | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI.                                            |

| oVirt UI Plugin                        |                  |          |
|----------------------------------------|------------------|----------|
| Platform                               | Platform Version | Comments |
| Red Hat Virtualization Manager (RHV-M) | 4.3, 4.4, 4.5    |          |

| VM File Level Recovery (Mountable Backups)           |           |                        |
|------------------------------------------------------|-----------|------------------------|
| Operating System                                     | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                        | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                        | Ubuntu    | XFS, EXT3, EXT4, BTRFS |

## Oracle Linux Virtualization Manager | OLVM

VM-00000001

### Oracle Linux Virtualization Manager (OLVM)

| Platform                                   | Platform Version    | Comments                                                                                     |
|--------------------------------------------|---------------------|----------------------------------------------------------------------------------------------|
| Oracle Linux Virtualization Manager (OLVM) | 4.2.8, 4.3.x, 4.4.x | Disk image transfer/ SSH transfer strategy – 4.3 +<br>Disk attachment strategy – 4.0 +       |
| Guest Tools (oVirt guest agent)            | Current             | Required only for VM FS freeze                                                               |
| Web browser                                | Google Chrome       | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |

### Oracle Linux Virtualization Manager UI Plugin

| Platform                                   | Platform Version | Comments |
|--------------------------------------------|------------------|----------|
| Oracle Linux Virtualization Manager (OLVM) | 4.3, 4.4         |          |

### VM File Level Recovery (Mountable Backups)

| Operating System                                     | OS Vendor | Feature Options        |
|------------------------------------------------------|-----------|------------------------|
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                        | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                        | Ubuntu    | XFS, EXT3, EXT4, BTRFS |

## Proxmox VE

| Proxmox VE       |                                                            |                                                                                              |
|------------------|------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| Platform         | Platform Version                                           | Comments                                                                                     |
| Proxmox VE       | 5.1, 5.2, 5.3, 5.4, 6.0, 6.1, 6.2, 6.3, 6.4, 7.0, 7.1, 7.2 |                                                                                              |
| QEMU guest agent | Current                                                    | Required only for VM FS freeze                                                               |
| Web browser      | Google Chrome                                              | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |

| VM File Level Recovery (Mountable Backups)           |           |                        |
|------------------------------------------------------|-----------|------------------------|
| Operating System                                     | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                        | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                        | Ubuntu    | XFS, EXT3, EXT4, BTRFS |

## KVM/Xen Legacy

| KVM/Xen Legacy   |                  |                                                                                                                         |
|------------------|------------------|-------------------------------------------------------------------------------------------------------------------------|
| Platform         | Platform Version | Comments                                                                                                                |
| KVM/Xen libvirt  | QEMU 2.1+        | VMs need libvirt with block commit feature using disks in the following formats - QCOW2, LVM, Ceph. Min. CentOS/RHEL 7. |
| QEMU guest agent | Current          | Required only for VM FS freeze                                                                                          |
| Web browser      | Google Chrome    | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI.                            |

| VM File Level Recovery (Mountable Backups)           |           |                        |
|------------------------------------------------------|-----------|------------------------|
| Operating System                                     | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                        | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                        | Ubuntu    | XFS, EXT3, EXT4, BTRFS |

## OpenStack Community | Vanilla

| OpenStack Community   Vanilla                                                                                                                                                                                                                               |                                                                         |                                                                                             |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| Platform                                                                                                                                                                                                                                                    | Platform Version                                                        | Comment                                                                                     |
| OpenStack Community                                                                                                                                                                                                                                         | Queens, Rocky, Stein, Train, Ussuri, Victoria, Wallaby, Xena, Yoga, Zed |                                                                                             |
| Source storage type supported                                                                                                                                                                                                                               | LVM, Ceph3, Ceph 4.x*, FC, iSCSI, NFS                                   | Ceph storage details in the notes below.                                                    |
| QEMU guest agent                                                                                                                                                                                                                                            | Current                                                                 | Required only for VM FS freeze                                                              |
| Web browser                                                                                                                                                                                                                                                 | Google Chrome                                                           | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI |
| <b>Notes:</b> <ul style="list-style-type: none"> <li>• Ceph3 = Red Hat Storage 3/Ceph 12.2+/luminous, CephFS, Ceph NBD</li> <li>• Ceph4 = Red Hat Storage 4/Ceph 14.2+/nautilus, CephFS, Ceph RBD</li> <li>• *Only Supported for OpenStack Train</li> </ul> |                                                                         |                                                                                             |

| OpenStack Horizon UI Plugin |                                                   |          |
|-----------------------------|---------------------------------------------------|----------|
| Platform                    | Platform Version                                  | Comments |
| OpenStack Community         | Train, Ussuri, Victoria, Wallaby, Xena, Yoga, Zed |          |

| VM File Level Recovery (Mountable Backups)           |           |                        |
|------------------------------------------------------|-----------|------------------------|
| Operating System                                     | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                        | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                        | Ubuntu    | XFS, EXT3, EXT4, BTRFS |

## Red Hat OpenStack

| Red Hat OpenStack Platform                                                                                                                                                                                                                                          |                                       |                                                                                              |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------|----------------------------------------------------------------------------------------------|
| Platform                                                                                                                                                                                                                                                            | Platform Version                      | Comment                                                                                      |
| Red Hat OpenStack Platform                                                                                                                                                                                                                                          | 13, 14, 15, 16, 16.1, 16.2, 17.0      |                                                                                              |
| Source storage type supported                                                                                                                                                                                                                                       | LVM, Ceph3, Ceph 4.x*, FC, iSCSI, NFS | Ceph storage details in the notes below.                                                     |
| QEMU guest agent                                                                                                                                                                                                                                                    | Current                               | Required only for VM FS freeze                                                               |
| Web browser                                                                                                                                                                                                                                                         | Google Chrome                         | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |
| Notes: <ul style="list-style-type: none"> <li>• Ceph3 = Red Hat Storage 3/Ceph 12.2+/luminous, CephFS, Ceph NBD</li> <li>• Ceph4 = Red Hat Storage 4/Ceph 14.2+/nautilus, CephFS, Ceph RBD</li> <li>• *Only Supported for Red Hat OpenStack Platform 16+</li> </ul> |                                       |                                                                                              |

| Red Hat OpenStack Platform Horizon UI Plugin |                      |          |
|----------------------------------------------|----------------------|----------|
| Platform                                     | Platform Version     | Comments |
| Red Hat OpenStack                            | 16, 16.1, 16.2, 17.0 |          |

| VM File Level Recovery (Mountable Backups)           |           |                        |
|------------------------------------------------------|-----------|------------------------|
| Operating System                                     | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                        | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                        | Ubuntu    | XFS, EXT3, EXT4, BTRFS |



## RDO OpenStack

| RDO OpenStack Platform                                                                                                                                                                                                                                   |                                                      |                                                                                              |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------|----------------------------------------------------------------------------------------------|
| Platform                                                                                                                                                                                                                                                 | Platform Version                                     | Comment                                                                                      |
| RDO OpenStack                                                                                                                                                                                                                                            | Queens, Rocky, Stein, Train, Victoria, Wallaby, Xena |                                                                                              |
| Source storage type supported                                                                                                                                                                                                                            | LVM, Ceph3, Ceph 4.x*, FC, iSCSI, NFS                | Ceph storage details in the notes below.                                                     |
| QEMU guest agent                                                                                                                                                                                                                                         | Current                                              | Required only for VM FS freeze                                                               |
| Web browser                                                                                                                                                                                                                                              | Google Chrome                                        | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |
| Notes: <ul style="list-style-type: none"> <li>• Ceph3 = Red Hat Storage 3/Ceph 12.2+/luminous, CephFS, Ceph NBD</li> <li>• Ceph4 = Red Hat Storage 4/Ceph 14.2+/nautilus, CephFS, Ceph RBD</li> <li>• *Only Supported for RDO OpenStack Train</li> </ul> |                                                      |                                                                                              |

| RDO OpenStack Horizon UI Plugin |                  |          |
|---------------------------------|------------------|----------|
| Platform                        | Platform Version | Comments |
| RDO OpenStack                   | Train            |          |

| VM File Level Recovery (Mountable Backups)           |           |                        |
|------------------------------------------------------|-----------|------------------------|
| Operating System                                     | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                        | Ubuntu    | XFS, EXT3, EXT4, BTRFS |

## Ubuntu OpenStack

| Ubuntu OpenStack                                                                                                                                                                                                                                                    |                                                                                                                               |                                                                                              |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| Platform                                                                                                                                                                                                                                                            | Platform Version                                                                                                              | Comment                                                                                      |
| Ubuntu OpenStack                                                                                                                                                                                                                                                    | 18.04 with Queens, Rocky, Stein, Train, Ussuri<br><br>20.04 with Ussuri, Victoria, Wallaby, Xena, Yoga<br><br>22.04 with Yoga | All distributions are with prefix Ubuntu +                                                   |
| Source storage type supported                                                                                                                                                                                                                                       | LVM, Ceph3, Ceph 4.x*, FC, iSCSI, NFS                                                                                         | Ceph storage details in the notes below.                                                     |
| QEMU guest agent                                                                                                                                                                                                                                                    | Current                                                                                                                       | Required only for VM FS freeze                                                               |
| Web browser                                                                                                                                                                                                                                                         | Google Chrome                                                                                                                 | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |
| Notes: <ul style="list-style-type: none"> <li>• Ceph3 = Red Hat Storage 3/Ceph 12.2+/luminous, CephFS, Ceph NBD</li> <li>• Ceph4 = Red Hat Storage 4/Ceph 14.2+/nautilus, CephFS, Ceph RBD</li> <li>• *Only Supported for Ubuntu 18.04 + OpenStack Train</li> </ul> |                                                                                                                               |                                                                                              |

| Ubuntu OpenStack Horizon UI Plugin |                                              |          |
|------------------------------------|----------------------------------------------|----------|
| Platform                           | Platform Version                             | Comments |
| Ubuntu OpenStack                   | Train, Ussuri, Victoria, Wallaby, Xena, Yoga |          |

| VM File Level Recovery (Mountable Backups)           |           |                        |
|------------------------------------------------------|-----------|------------------------|
| Operating System                                     | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                        | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |

|                               |        |                        |
|-------------------------------|--------|------------------------|
| Ubuntu 14.x, 16.x, 18.x, 20.x | Ubuntu | XFS, EXT3, EXT4, BTRFS |
|-------------------------------|--------|------------------------|

### Virtuozzo

| Virtuozzo        |                    |                                                                                              |
|------------------|--------------------|----------------------------------------------------------------------------------------------|
| Platform         | Platform Version   | Comment                                                                                      |
| Virtuozzo        | 4.7, 5.0, 5.1, 5.2 | All distributions are with prefix Ubuntu +                                                   |
| QEMU guest agent | Current            | Required only for VM FS freeze                                                               |
| Web browser      | Google Chrome      | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |
| Notes:           |                    |                                                                                              |

| VM File Level Recovery (Mountable Backups)           |           |                        |
|------------------------------------------------------|-----------|------------------------|
| Operating System                                     | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                        | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                        | Ubuntu    | XFS, EXT3, EXT4, BTRFS |

## Oracle VM

| Oracle VM   |                  |                                                                                              |
|-------------|------------------|----------------------------------------------------------------------------------------------|
| Platform    | Platform Version | Comments                                                                                     |
| Oracle VM   | 3.4 +            | Disk image transfer/ SSH transfer strategy – 4.2 +<br>Disk attachment strategy – 4.0 +       |
| Web browser | Google Chrome    | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |

| VM File Level Recovery (Mountable Backups)           |           |                        |
|------------------------------------------------------|-----------|------------------------|
| Operating System                                     | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                        | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                        | Ubuntu    | XFS, EXT3, EXT4, BTRFS |

## Citrix Hypervisor (formely Xen Server)

| Citrix Hypervisor   XenServer                                                                                                                                                                                        |                                                       |                                                                                              |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------|----------------------------------------------------------------------------------------------|
| Platform                                                                                                                                                                                                             | Platform Version                                      | Comments                                                                                     |
| Citrix Hypervisor                                                                                                                                                                                                    | 6.5, 7.0, 7.1, 7.2, 7.3, 7.4, 7.5, 7.6, 8.0, 8.1, 8.2 |                                                                                              |
| Citrix VM Tools                                                                                                                                                                                                      | Current                                               | Required only for VM FS freeze                                                               |
| Web browser                                                                                                                                                                                                          | Google Chrome                                         | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |
| <b>Notes:</b> <ul style="list-style-type: none"> <li>Change-block Tracking strategy available from 7.3 +</li> <li>Changed block tracking is available only to customers with XenServer Enterprise Edition</li> </ul> |                                                       |                                                                                              |

| VM File Level Recovery (Mountable Backups)                                                                      |           |                        |
|-----------------------------------------------------------------------------------------------------------------|-----------|------------------------|
| Operating System                                                                                                | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                                                                              | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                                                                                        | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                                                                                      | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                                                                                      | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                                                                                   | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x                                                                        | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x                                                            | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                                                                             | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                                                                           | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                                                                                   | Ubuntu    | XFS, EXT3, EXT4, BTRFS |
| <b>Notes:</b> <ul style="list-style-type: none"> <li>Only applied to Change-Block Tracking strategy.</li> </ul> |           |                        |

## XCP-ng

Version 7.4.0

| XCP-ng             |                              |                                                                                              |
|--------------------|------------------------------|----------------------------------------------------------------------------------------------|
| Platform           | Platform Version             | Comments                                                                                     |
| XCP-ng             | 7.4, 7.5, 7.6, 8.0, 8.1, 8.2 |                                                                                              |
| XCP-ng guest tools | Current                      | Required only for VM FS freeze                                                               |
| Web browser        | Google Chrome                | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |

| VM File Level Recovery (Mountable Backups)                                                       |           |                        |
|--------------------------------------------------------------------------------------------------|-----------|------------------------|
| Operating System                                                                                 | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                                                               | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                                                                         | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                                                                       | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                                                                       | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                                                                    | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x                                                         | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x                                             | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                                                              | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                                                            | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                                                                    | Ubuntu    | XFS, EXT3, EXT4, BTRFS |
| Notes:                                                                                           |           |                        |
| <ul style="list-style-type: none"> <li>Only applied to Change-Block Tracking strategy</li> </ul> |           |                        |

## Microsoft Hyper-V

| Hyper-V     |                  |                                                                                              |
|-------------|------------------|----------------------------------------------------------------------------------------------|
| Platform    | Platform Version | Comments                                                                                     |
| Hyper-V     | 2016, 2019, 2022 | Supported virtual machine configuration versions 6.2+                                        |
| Hyper-V VMM | Current          | Required only for VM FS freeze                                                               |
| Web browser | Google Chrome    | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |

| VM File Level Recovery (Mountable Backups)           |           |                        |
|------------------------------------------------------|-----------|------------------------|
| Operating System                                     | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                        | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                        | Ubuntu    | XFS, EXT3, EXT4, BTRFS |

## Azure Stack HCI

| Hyper-V         |                  |                                                                                              |
|-----------------|------------------|----------------------------------------------------------------------------------------------|
| Platform        | Platform Version | Comments                                                                                     |
| Azure Stack HCI | 21H2, 22H2       | Supported virtual machine configuration versions 6.2+                                        |
| Web browser     | Google Chrome    | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |

| VM File Level Recovery (Mountable Backups)           |           |                        |
|------------------------------------------------------|-----------|------------------------|
| Operating System                                     | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                        | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                        | Ubuntu    | XFS, EXT3, EXT4, BTRFS |



## Huawei FusionCompute

| FusionCompute             |                  |                                                                                              |
|---------------------------|------------------|----------------------------------------------------------------------------------------------|
| Platform                  | Platform Version | Comments                                                                                     |
| FusionCompute             | 8.0+             |                                                                                              |
| FusionCompute guest tools | Current          | Required for incremental backups                                                             |
| Web browser               | Google Chrome    | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |

| VM File Level Recovery (Mountable Backups)           |           |                        |
|------------------------------------------------------|-----------|------------------------|
| Operating System                                     | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                        | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                        | Ubuntu    | XFS, EXT3, EXT4, BTRFS |

## Scale Computing HC3

| Scale Computing HC3             |                  |                                                                                              |
|---------------------------------|------------------|----------------------------------------------------------------------------------------------|
| Platform                        | Platform Version | Comments                                                                                     |
| Scale Computing HC3             | 8.9+             |                                                                                              |
| Scale Computing HC3 guest tools | Current          | Required only for VM FS freeze                                                               |
| Web browser                     | Google Chrome    | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |

| VM File Level Recovery (Mountable Backups)           |           |                        |
|------------------------------------------------------|-----------|------------------------|
| Operating System                                     | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                        | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                        | Ubuntu    | XFS, EXT3, EXT4, BTRFS |

## OpenNebula

| OpenNebula       |                  |                                                                                              |
|------------------|------------------|----------------------------------------------------------------------------------------------|
| Platform         | Platform Version | Comments                                                                                     |
| OpenNebula       | 6.6+             |                                                                                              |
| QEMU guest agent | Current          | Required only for VM FS freeze                                                               |
| Web browser      | Google Chrome    | The latest version of the Google Chrome browser in order to access the Dell EMC vProtect UI. |

| VM File Level Recovery (Mountable Backups)           |           |                        |
|------------------------------------------------------|-----------|------------------------|
| Operating System                                     | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                        | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                        | Ubuntu    | XFS, EXT3, EXT4, BTRFS |

## Supported Platforms | Storage Providers

### Filesystem

Filesystem backup source

| Operating System                         | OS Vendor | Comments                                                    |
|------------------------------------------|-----------|-------------------------------------------------------------|
| CentOS 7.x, 8.x, Stream                  | CentOS    | Any POSIX compliant filesystem mounted on the vProtect Node |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x | Red Hat   | Any POSIX compliant filesystem mounted on the vProtect Node |

#### Storage File Level Recovery (Mountable Backups)

| Operating System                         | OS Vendor | Feature Options |
|------------------------------------------|-----------|-----------------|
| CentOS 7.x, 8.x, Stream                  | CentOS    | -               |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x | Red Hat   | -               |

## Ceph RBD

| Platform | Platform Version                       | Comments             |
|----------|----------------------------------------|----------------------|
| CEPH     | 13.2.x, 14.2.x, 15.2.x, 16.2.x, 17.2.x | Only RBD-NBD volumes |

| Storage Ceph RBD Recovery (Mountable Backups)        |           |                        |
|------------------------------------------------------|-----------|------------------------|
| Operating System                                     | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                   | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                             | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                           | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                        | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x             | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x, 12.x, 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                  | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x, 16.x, 18.x, 20.x                        | Ubuntu    | XFS, EXT3, EXT4, BTRFS |

## Nutanix Files

| Platform                   | Platform Version                                                                                                                                                                                                                | Comments |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| AOS, Nutanix Files Version | 3.5.1, 3.5.2, 3.5.3, 3.5.4, 3.5.5,<br>3.5.6, 3.6.0, 3.6.1, 3.6.1.1,<br>3.6.1.2, 3.6.1.3, 3.6.2, 3.6.3,<br>3.6.4, 3.6.5, 3.7.0, 3.7.1, 3.7.2,<br>3.7.2.1, 3.7.3, 3.8.0, 3.8.0.1,<br>4.0.0.2, 4.1.0.1, 4.1.0.2,<br>4.1.0.3, 4.2.0 |          |

### Storage File Level Recovery (Mountable Backups)

| Protocol | OS Vendor | Feature Options |
|----------|-----------|-----------------|
| NFS, SMB | -         | -               |

## Nutanix Volume Groups

| Platform           | Platform Version                                | Comments |
|--------------------|-------------------------------------------------|----------|
| AOS, Nutanix Prism | 5.11- 5.15 LTS, 5.17, 5.18, 5.19, 5.20 LTS, 6.0 |          |

| Nutanix Volume Groups Recovery (Mountable Backups)     |           |                        |
|--------------------------------------------------------|-----------|------------------------|
| Operating System                                       | OS Vendor | Feature Options        |
| Windows 8, 8.1, 10                                     | Microsoft | NTFS, FAT32            |
| Windows 2012, 2016, 2019                               | Microsoft | NTFS, FAT32            |
| CentOS 7.x                                             | CentOS    | XFS, EXT3, EXT4        |
| CentOS 8.x                                             | CentOS    | XFS, EXT3, EXT4        |
| CentOS Stream                                          | CentOS    | XFS, EXT3, EXT4        |
| Red Hat Enterprise Linux (RHEL) 7.x, 8.x               | Red Hat   | XFS, EXT3, EXT4        |
| SUSE Linux Enterprise Server (SLES) 11.x , 12.x., 15.x | SUSE      | XFS, EXT3, EXT4, BTRFS |
| OpenSUSE 13.x, 15.x                                    | SUSE      | XFS, EXT3, EXT4, BTRFS |
| Oracle Linux 7.x, 8.x                                  | Oracle    | XFS, EXT3, EXT4        |
| Ubuntu 14.x., 16.x, 18.x, 20.x                         | Ubuntu    | XFS, EXT3, EXT4, BTRFS |

**DECISION**

SERVIÇOS DE TECNOLOGIA DA INFORMAÇÃO Ltda.

# **ATESTADOS** **TÉCNICOS**

DECISION

**Brasília (Sede)**

Setor Hoteleiro Sul - Quadra 08 - Conjunto 31  
Bloco A - Sala 102 - Asa Sul - Brasília/DF  
Cap. 70.322-915 - Tel. (61) 3045.0050

**Salvador**

Avenida Terceiro Neves, 620 - Salas 2910 e 2911  
29º andar - Torre Empresarial do Ed. Mundo Plaza  
Caminho das Árvores - Salvador/BA - Cap. 41.820-500  
Tel. (71) 3265.7007

**São Paulo**

Rua Arizona, 1.422 - Conjunto 06 - Ed. Platinum  
Building Berrini - Berrini - São Paulo/SP - Cap. 04.567-003  
Tel. (11) 5583.0344



## ATESTADO DE CAPACIDADE TÉCNICA

Atestamos, para os devidos fins, que a empresa DECISION Serviços de Tecnologia da Informação Ltda, inscrita no CNPJ nº 03.535.902/0001-10, situada no SHS Qd. 06 Conjunto A Bloco A sala 807 Parte A – Brasília- DF, forneceu, instalou e configurou, bem como apoiou tecnicamente os produtos listados abaixo, de propriedade da Agência Nacional de Aviação Civil, inscrita no CNPJ sob nº 07.947.821/0001-89, situada no Edifício Parque Cidade Corporate, Setor Comercial Sul, Quadra 09, Lote C, Torre A, 1º andar, Brasília – DF, conforme contrato nº 16/2014.


| Qtd. | Descrição do Produto          |
|------|-------------------------------|
| 2    | EMC. Data Domain DD 4200      |
| 2    | EMC. Backup Suite - Networker |

EM complemento aos serviços prestados, foram realizados Backup em 48 (quarenta e oito) servidores físicos, bem como em ambiente virtual clusterizado com Microsoft Hyper-V, composto de 15 (quinze) Hosts Físicos.

Declaramos que os serviços foram executados com qualidade satisfatória, cumprindo plenamente as exigências contratuais e que não constam em nossos arquivos fatos a repostar que desabonem a conduta desta empresa.

Atenciosamente.

Brasília, 10 de dezembro de 2019.



\_\_\_\_\_  
*Marcelo Nogueira Lino*  
Gerente de Infraestrutura Tecnológica  
Marcelo Nogueira Lino INSC: 2126657  
Gerência de Infraestrutura Tecnológica - GEIT

24

24º OFÍCIO DE NOTAS - José Mario Pinheiro Pinto

Avenida Almirante Balthus, 139 - C - Centro - Rio de Janeiro - RJ 20047-900

089807A0104120

**AUTENTICAÇÃO**

Certifico e dou fé que a presente cópia é a reprodução  
do documento que me foi apresentado como sendo o original.

Autenticação: R\$ 2,46

Proc. dados: R\$ 6,00

Total: R\$ 8,46

Selo: **EDKM73542-DSJ**

Consulte em <https://www3.tjrj.tju.br/sitepublico>



AUTENTICAÇÃO

## ATESTADO DE CAPACIDADE TÉCNICA

Atestamos para os devidos fins que a empresa **DECISION SERVIÇOS DE TECNOLOGIA DA INFORMAÇÃO LTDA**, inscrita no CNPJ N° 03.535.902/0001-10, situada no ST SETOR HOTELEIRO SUL QUADRA 06, CONJ. A, BLOCO A, SALA 807 – ASA SUL – Brasília/DF, forneceu para a **COMPANHIA DE GÁS DO ESTADO DA BAHIA – BAHAGÁS**, inscrita no CNPJ N° 34.432.153/0001-20, situada na Av. Tancredo Neves, 450, Ed. Suarez Trade, 20° andar, Caminho das Árvores, Salvador/BA, os seguintes equipamentos e *softwares*:

| DESCRIÇÃO                                                                         | QUANT. |
|-----------------------------------------------------------------------------------|--------|
| Software de Backup e Restore Simpana da Commvault                                 | 01     |
| Solução de backup em disco com deduplicação EMC Data Domain DD620                 | 02     |
| Fitoteca automatizada Tape Library Quantum Scalar 180                             | 01     |
| Fita de backup LTO-6 Quantum Data Cartridge com etiqueta de código de barras      | 50     |
| Fita de limpeza LTO-6 Quantum Cleaning Cartridge com etiqueta de código de barras | 01     |

Todos os equipamentos e *softwares* foram fornecidos com garantia do fabricante de 3 (três) anos. Declaramos que os equipamentos e *softwares* foram entregues dentro do prazo contratado, com excelente qualidade técnica, nada tendo a reportar que desabone a Empresa **DECISION SERVIÇOS DE TECNOLOGIA DA INFORMAÇÃO LTDA**.

Salvador, 12 de Junho de 2014.



*Patricia Zarife Cox*

**Patrícia Zarife Cox**

Gerente de Tecnologia da Informação

**Patrícia Zarife Cox**

Gerente de Tecnologia  
da Informação

24

24º OFÍCIO DE NOTAS - José Mario Pinheiro Pinto 089607A/2596583  
Avenida Almirante Barros, 139 - C - Centro - Rio de Janeiro - RJ - CEP: 20121-000 - Fone: (21)3551-0021

**AUTENTICAÇÃO**

Certifico e dou fé que o presente documento que me foi apresentado como sendo o original, é autêntico e verdadeiro.  
Autenticação: R\$ 2,35  
Proc. dado: R\$ 5,78  
Total: R\$ 8,13

Selo: EDGA51878-MDM

Consulte em <http://www3.tjrr.jus.br/ajrepública>



AUTENTICAÇÃO

## ATESTADO DE CAPACIDADE TÉCNICA

A CAESB - Companhia de Saneamento Ambiental do Distrito Federal, situada à Av.Sibipiruna lotes 13/21 Aguas Claras Brasília - DF **atesta** para todos os fins que a empresa **DECISION Serviços de Tecnologia da Informação Ltda**; com sede na SHS QD 06 Conjunto A Bloco F SL 807 - Brasil 21 - Asa Sul Brasília DF - CEP 70322-915 inscrita no CNPJ/MF nº 03.535.902/0001-10 referente ao pregão **PE Nº 90/2012, processo 092.002582/2012**, contrato nº **8261** assinado na data de 14/08/2012 forneceu sistemas de armazenamento de dados unificados e dispositivos de backup em disco, incluindo instalação, configuração, treinamento e suporte conforme descrição abaixo:

| Qtd. | Descrição                                                                                               | Marca       |
|------|---------------------------------------------------------------------------------------------------------|-------------|
| 2    | Solução de armazenamento de dados, composta Storages VNX5300 com garantia de 36 meses                   | EMC         |
| 2    | Switch FibreChannel, com 24 portas DS5100B com garantia de 36 meses                                     | EMC/Brocade |
| 2    | Solução de backup em disco com 8TB liquido Modelo : DataDomain DD640 com 1 x Chassis e 12 x Discos 1TB. | EMC         |
| 1    | Switch FibreChannel, com 24 portas DS5100B                                                              | EMC/Brocade |
| 1    | Instalação e implementação da solução de armazenamento e conectividade (Switch)                         | Decision    |
| 1    | Serviços de Migração de Dados                                                                           | Decision    |
| 1    | Serviços de Treinamento                                                                                 | EMC         |
| 1    | Serviços de Garantia e Assistência Técnica por 36 meses (Storages e Switch)                             | EMC         |

Por ser verdade, firmamos o presente atestado,

Brasília 16 de Novembro de 2012



---

**Marilson Manoel de Barros Santos**  
Caesb Companhia de Saneamento do Distrito Federal  
Coordenador de Infraestrutura de TI - PRTI  
Fone: 3213-7173 - 7140

**24** 24º OFÍCIO DE NOTAS - José Maria de Faria Pinto 0826074480859  
Avenida Alameda Santos, 133 - C. Centro - São Paulo - SP - CEP: 01311-001

**AUTENTICAÇÃO**  
Certifico e dou fé que a presente cópia é reprodução fiel do documento que foi apresentado como sendo o original.

Rio de Janeiro, 1 de setembro de 2023.

**24** OFÍCIO DE NOTAS  
MST. SAJ/RS51  
Tel. 2508-8000



JARA LIMA RUAS PAIVA PINTO  
Emolumentos: R\$ 7,41 - T.J. Fundos: R\$ 3,02 + R\$ 1,50 = R\$ 11,91  
Selo: EEOY75133.AMG - Consulte em <http://www4.tj.jus.br/cital/estrutural/consultaselo/>

## ATESTADO DE CAPACIDADE TÉCNICA

O Centro Brasileiro de Pesquisa em Avaliação e Seleção e de Promoção de Eventos (Cebraspe), associação civil sem fins lucrativos, qualificada como Organização Social (OS) pela Presidência da República, por meio do Decreto n.º 8.078, de 19 de agosto de 2013, com sede no campus Universitário Darcy Ribeiro, Asa Norte, em Brasília/DF, inscrito no CNPJ n.º 18.284.407/0001-53, registrado no 2.º Ofício de Registros de Pessoas Jurídicas de Brasília/DF, sob o n.º 000082415, em 13 de maio de 2013, **atesta**, para fins de comprovação de capacidade técnica, que a empresa **Decision Serviços de Tecnologia da Informação Ltda**, inscrita no CNPJ n.º 03.535.902/0001-10, sediada na SHS quadra 06, conjunto A, bloco A, sala 08, Ed. Brasil 21, Brasília/DF, forneceu para este Centro os produtos e serviços de instalação, migração e configuração, referente à aquisição de equipamentos de armazenamento digital com garantia do fabricante, objeto do Pregão Eletrônico n.º 12/2015, Ordem de Fornecimento n.º 2000000531 e Processo Administrativo n.º 1178/2015, em conformidade com os dados abaixo especificados:

| Item | Descrição                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Unid. | Quant. | Valor R\$    |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------|--------|--------------|
| 1    | <p><b>Armazenamento tipo NAS expansível – ISILON</b></p> <p>Especificação Técnica Mínima:</p> <p>1.1.Ser 1 (um) nó de ISILON da Série X 410 com armazenamento, processamento, memória e conectividade;</p> <p>1.2.Com 64 (sessenta e quatro) Gb de memória bruta para cache;</p> <p>1.3.Com 2 (dois) processadores XEON octa-core;</p> <p>1.4.Com duas interfaces 1 GBPS ethernet em RJ45 e duas interfaces 10 GBPS com módulos SFP+ (10 BASE-SR);</p> <p>1.5.Com 2 cabos para conexão de back end em infiniband;</p> <p>1.6.Com 6 (seis) discos SSD de 800 GB e 30 (trinta) discos SATA de 4T em um mesmo nó para um total combinado de 122,4 TB bruto, garantindo performance por uso da tecnologia SSD;</p> <p>1.7.Deve se integrar ao cluster Isilon existente, formando um novo Disk Pool, ampliando a capacidade do cluster existente;</p> <p>1.8.Licenciado para uso de:</p> <p>a. SMB</p> <p>b. Smart Pools</p> <p>c. Smart Coonect</p> <p>d. SnapShotIQ</p> <p>e. InsightIQ</p> <p>f. SmartLock – Enterprise</p> | Un    | 01     | 1.049.000,00 |

|   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |    |    |            |
|---|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|----|------------|
|   | g. SmartDedup – High Density<br>h. Armazenamento de Objeto<br>1.9.Com todos os cabos e acessórios necessários ao funcionamento pleno;<br>1.10. Com os serviços de instalação, configuração, operação assistida, consultoria, suporte técnico e migração dos dados.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |    |    |            |
| 2 | <b>DataDomain</b><br>O equipamento deverá no mínimo:<br>3.1.Ser um DataDomain da família 2.500;<br>3.2.Possuir área de armazenamento disponível a backup de 30 TB;<br>3.3.Possuir duas interfaces SFP+ para comunicação de rede, com os conversores ópticos no padrão LC Multimodo;<br>3.4.Possuir quatro interfaces Ethernet RJ 45 de 1Gbps;<br>3.5.Com licenciamento para uso de EMC Data Domain Boost; para até 500 usuários;<br>3.6.Com licenciamento para uso de EMC Data Domain Encryption para até 500 usuários; 3.7.Com licenciamento para uso de EMC Data Domain Management Solutions;<br>3.8.Com todos os cabos e acessórios necessários ao seu funcionamento;<br>3.9.Com os serviços do fabricante para instalação, configuração e migração dos dados. | Un | 01 | 590.000,00 |

Atesta-se, ainda, que tais produtos e serviços foram fornecidos/prestados satisfatoriamente, não existindo registros, até a presente data, de fatos que desabonem a conduta e a responsabilidade com as obrigações assumidas pela referida empresa.

Brasília/DF, 18 de março de 2019.

Adriana Rigon Weska  
 Diretora-Geral



## ASSINATURA(S) ELETRÔNICA(S)



A autenticidade do documento pode ser conferida no site:  
<http://ged-01/ValidarDocumento.aspx>  
informando o código CRC: 41757642646B44593942633D / Página 3 de 3



Assinado digitalmente por: ADRIANA RIGON WESKA, Certificado Digital: CN=ADRIANA RIGON  
WESKA:34691723153, OU=AR NATHYELLE, OU=VALID, OU=RFB e-CPF A1, OU=Secretaria da Receita Federal do  
Brasil - RFB, C=ICP-Brasil, C=BR Data da Assinatura: 20/03/2019 12:37:04

## ATESTADO DE CAPACIDADE TÉCNICA

Atestamos para os devidos fins que a empresa **DECISION SERVIÇOS DE TECNOLOGIA DA INFORMAÇÃO LTDA.**, inscrita no CNPJ sob o nº **03.535.902/0001-10**, situada a SHS Quadra 06, Conjunto A, Bloco A, Sala 807 – Brasília / DF, forneceu de forma satisfatória os itens descritos abaixo, para a **STATOIL BRASIL OLÉO E GÁS**, inscrito no CNPJ nº **04.028.583/0001-10**, situado à Rua do Russel, 804 – 6º/7º andar – Salas 601, 602, 701 e 702 - Rio de Janeiro / RJ, incluindo serviços de instalação e assistência técnica pelo período de 36 (trinta e seis) meses com suporte on-site 24x7, os seguintes equipamentos.

|   |                                                           |                                     |   |
|---|-----------------------------------------------------------|-------------------------------------|---|
| 1 | EMC DataDomain DD-160 com 4Tb de proteção de dados        | 36 meses de garantia em regime 24x7 | 1 |
| 2 | Licença de replicação de backup EMC DataDomain Replicator | 36 meses de garantia em regime 24x7 | 1 |
| 3 | Licença de VTL/NDMP para EMC DataDomain                   | 36 meses de garantia em regime 24x7 | 1 |
| 4 | Interfaces 1Gbps Dual Cooper                              | 36 meses de garantia em regime 8x5  | 2 |

Declaramos que não constam em nossos arquivos fatos a reportar que desabonem a conduta da referida empresa.

Rio de Janeiro, 21 de Agosto de 2012.

  
Breno de Moura Spagnuolo Gomes

Coordenador de TI

Tel.: +55 21 3479-9800

E-mail: bsg@statoil.com

---

**Endereço**

Rua do Russel nº 804  
22210-010 - Rio de Janeiro, RJ  
Brasil

**Telefone**

+ 55 (21) 3479-9800

**Telefax**

+ 55 (21) 3479-9899

24

24º OFÍCIO DE NOTAS - José Mario Pinheiro Pinto (092074-080137)  
Avenida Almeida Barros, 129 - C. Centro - Rio de Janeiro - RJ - CEP: 21003-601 - Telefons: 210053-6021

**AUTENTICAÇÃO**  
Certifico e dou fé que a presente cópia é a reprodução fiel do documento que me foi apresentado como sendo original.  
Rio de Janeiro, 27 de setembro de 2021.



DANIEL DE BARROS MONTEIRO

Emolumentos: R\$ 8,25 - T.J.+Furto(s): R\$ 2,54 - Total: R\$ 8,8

Selo: EDYG39754-ANG - Consulte em <http://www3.trj.jus.br/sitapublico>

AUTENTICAÇÃO



Tribunal de Contas dos Municípios  
do Estado da Bahia

## ATESTADO DE CAPACIDADE TÉCNICA

Atestamos para os devidos fins que a empresa **DECISION SERVIÇOS DE TECNOLOGIA DA INFORMAÇÃO LTDA**, inscrita no CNPJ sob o nº **03.535.902/0002-00**, estabelecida na Avenida Tancredo Neves, 620, Edifício Mundo Plaza Empresarial 29º andar Sala 2910/2911, Caminho das Árvores, Salvador – BA, Cep: 41.820-020, forneceu para o **Tribunal de Contas dos Municípios do Estado da Bahia**, inscrito no CNPJ sob o nº 32.634.420/0001-16, com sede à Av. IV, nº 495, 3º andar, Centro Administrativo da Bahia, nesta Capital, Solução de Hiperconvergência, conforme descrito abaixo, incluindo serviços de implementação do ambiente VMware, instalação, configuração, atualizações, treinamento hands-on da solução e suporte técnico on-site 24x7, 60 meses.

**EDITAL:** PREGÃO ELETRÔNICO Nº 008/2017

**CONTRATO NO.007/2018**

**VIGÊNCIA CONTRATUAL:** 13/08/2018 à 13/08/2019

**GARANTIA:** 12 (doze) meses

| Qtd | Discriminação                |
|-----|------------------------------|
| 4   | VxRail P570F                 |
| 4   | Data Protection Suite Bundle |
| 2   | DataDomain DD3300            |
| 4   | Dell Networking S4048        |
| 1   | Dell Rack Netshelter SX 42U  |

**Gerente de Projeto:** André Reis

**Técnicos Implementadores:** Pedro Azevedo

Ressaltamos que a solução foi entregue e implementada com garantia e qualidade, e que os profissionais envolvidos demonstraram capacidade técnica e realizaram, de forma satisfatória, os serviços previstos em contrato.

A empresa vem sendo correta em suas relações comerciais e técnicas, demonstrando idoneidade moral no cumprimento do contrato, nada havendo a reportar que desabone a sua conduta

Salvador, 25 de setembro de 2019.

Nome: Pedro Vieira

DTI - Diretoria de Tecnologia da Informação

e-mail: [pedro.vieira@tcm.ba.gov.br](mailto:pedro.vieira@tcm.ba.gov.br)

Tel.: (71) 3115-5613

24

24º OFÍCIO DE NOTAS - José Mario Pinheiro Pinto 389687AG825744  
Avenida Afirante Brasil, 139 - C - Centro - Rio de Janeiro - RJ - CEP: 20053-902 - Fone: (21) 3553-6021

**AUTENTICAÇÃO**

Certifico e dou fé que o presente cópia e reprodução do documento que me foi apresentado com o número 0044

Autenticação: R\$ 2,46  
Proc. dedoe: R\$ 6,00

Total: R\$ 8,46

Selo: EDIH22213-VTN

Consulte em <http://www3.trf1.jus.br/sitepublico>



AUTENTICAÇÃO

Eldorado do Sul, 27 de setembro de 2023

À  
DEFENSORIA PÚBLICA DO ESTADO DO RIO DE JANEIRO  
A/C Sr. Pregoeiro

Ref: PE 027/2023 Processo: E-20/001.007805/2022

## DECLARAÇÃO TÉCNICA

**DELL COMPUTADORES DO BRASIL LTDA. ("Dell")**, inscrita no CNPJ/MF sob o nº 72.381.189/0001-10, com sede na Av. Industrial Belgraf, 400 – Medianeira – CEP 92990-000, Eldorado do Sul/RS, com o objetivo de complementar as informações que não constam no Catálogo Técnico Oficial do(s) produto(s) abaixo ofertado(s), vem, através da presente, declarar o que segue:

Objeto: PowerProtect DD6900

Declaramos que os equipamentos são novos, de primeiro uso e estão em linha de fabricação. Além de serem acondicionados em embalagem individual adequada, com o menor volume possível, utilizando materiais recicláveis, de forma a garantir a máxima proteção durante o transporte e o armazenamento.

DELL COMPUTADORES DO BRASIL  
LTD.A.72381189000110

Digitally signed by DELL  
COMPUTADORES DO BRASIL  
LTD.A.72381189000110  
Date: 2023.09.27 11:25:49 -0300

---

**Dell Computadores do Brasil Ltda**  
**Juliane Casagrande Rodrigues – Gerente de Vendas**

1

Eldorado do Sul, 28 de setembro de 2023.

À  
DEFENSORIA PÚBLICA DO ESTADO DO RIO DE JANEIRO  
Ref.: Processo E-20/001.007805/2022 – Pregão eletrônico 027/2023

## DECLARAÇÃO

A **Dell Computadores do Brasil Ltda.** ("Dell") inscrita no CNPJ sob o n. 72.381.189.0001-10 e com sede na Av. Industrial Belgraf n. 400, Eldorado do Sul, RS, vem, por meio de seu representante legal, declarar que a empresa **DECISION SERVIÇOS DE TECNOLOGIA DA INFORMAÇÃO LTDA**, com sede na, RUA: GILDASIO AMADO, Nº 55, SALA 1.604 CEP: 22.631-020 – BARRA DA TIJUCA – RJ inscrita no CNPJ sob o nº 03.535.902/0005-44, faz parte do Programa de Parceria DELL TECHNOLOGIES e é atualmente parceira da Dell, estando autorizada a comercializar os produtos e serviços de renovação de garantia Dell em todo o território brasileiro.

DELL COMPUTADORES DO BRASIL  
LTDA 72381189000110 Digitally signed by DELL  
COMPUTADORES DO BRASIL  
LTDA 72381189000110  
Date: 2023.09.28 08:43:38 -0300

**Dell Computadores do Brasil Ltda.**

**Juliane Casagrande Rodrigues – Gerente de Vendas**

**DELL Computadores do Brasil Ltda.**

Av. Industrial Belgraf, 400, Eldorado do Sul / RS. Gerat : 51 3481 5500 Fax : 51 3481 5458.

## FILTROS APLICADOS:

Busca livre: 03535902000544

Cadastro: CEIS

LIMPAR

Data da consulta: 21/09/2023 20:02:00

Data de última atualização: 01/1900 (Sistema Integrado de Registro do CEIS/CNEP - Acordos de Leniência), 09/2023 (Sistema Integrado de Administração Financeira do Governo Federal (SIAFI) - CEPIM), 09/2023 (Diário Oficial da União - CEAF), 09/2023 (Sistema Integrado de Registro do CEIS/CNEP - CEIS), 09/2023 (Sistema Integrado de Registro do CEIS/CNEP - CNEP)

| DETALHAR                    | CADASTRO | CNPJ/CPF SANCIONADO | NOME SANCIONADO | UF SANCIONADO | ÓRGÃO/ENTIDADE SANCIONADORA | CATEGORIA SANÇÃO | DATA DE PUBLICAÇÃO DA SANÇÃO | VALOR DA MULTA | QUANTIDADE |
|-----------------------------|----------|---------------------|-----------------|---------------|-----------------------------|------------------|------------------------------|----------------|------------|
| Nenhum registro encontrado. |          |                     |                 |               |                             |                  |                              |                |            |



**DECISION**

SERVIÇOS DE TECNOLOGIA DA INFORMAÇÃO S/A

# DECLARAÇÕES

DECISION

**Brasília (Sede)**  
Salar Holístico Sul - Quadra 06 - Conjunto W  
Bloco A - Sala 102 - Asa Sul - Brasília/DF  
Cep. 70.332-910 - Tel. (61) 3048.0050

**Salvador**  
Avenida Tancredi Neves - 820 - Sala 2910 e 2911  
29º andar - Torre Empresarial do Ed. Mundi Plaza  
Carmem das Águas - Salvador/BA - Cep. 41.825-020  
Tel. (71) 3565.7007

**São Paulo**  
Rua Arizona, 1.423 - Conjunto 76 - Ed. Platinum  
Building Berrini - Berrini - São Paulo/SP - Cep. 04.567-020  
Tel. (11) 6583.0344

## ANEXO IV

### DECLARAÇÃO DE ATENDIMENTO AO DISPOSTO NO ART. 7.º, INCISO XXXIII, DA CONSTITUIÇÃO FEDERAL

**DEFENSORIA PÚBLICA DO ESTADO DO RIO DE JANEIRO**  
**EDITAL DE PREGÃO ELETRÔNICO Nº 027/23**  
Processo nº E-20/001.007805/2022

DECLARO, sob as penas da Lei, em atendimento ao Edital do **Pregão Eletrônico nº PE 027/23**, promovido por essa DPRJ, na Avenida Marechal Câmara nº 314, Centro – Rio de Janeiro/RJ CEP: 20.020-080, que a empresa **DECISION SERVIÇOS DE TECNOLOGIA DA INFORMAÇÃO LTDA**, CNPJ nº **03.535.902/0005-44**, com sede à **RUA GILDASIO AMADO, 55 - SALA 1604 - BARRA DA TIJUCA - RIO DE JANEIRO/RJ - CEP: 22.631-020**, por mim representada, não possui em seu quadro funcional nenhum menor de 18 (dezoito) anos desempenhando trabalho noturno, perigoso ou insalubre ou qualquer trabalho por menor de 16 (dezesesseis) anos, em obediência ao art. 7º, inciso XXXIII, da Constituição Federal.

Rio de Janeiro (RJ), 28 de setembro de 2023.



**ALVARO LUIZ SOARES FUZEIRO**  
**GERENTE DE CONTAS**  
alvaro.fuzeiro@decision-tec.com.br  
Tel. (21) 99649-5376  
ID 05.568.662-0 (IFP/RJ)  
CPF 839.674.197-20

**03.535.902/0005-44**  
DECISION SERVIÇOS DE  
TECNOLOGIA DA INFORMAÇÃO LTDA  
Rua Gildasio Amado 00055 Sala 1604  
Barra Da Tijuca - Cep: 22631-020  
RIO DE JANEIRO RJ

**ANEXO V**  
**DECLARAÇÃO EM ATENDIMENTO À LEI 7.258/2016**

**DEFENSORIA PÚBLICA DO ESTADO DO RIO DE JANEIRO**  
**EDITAL DE PREGÃO ELETRÔNICO Nº 027/23**  
Processo nº E-20/001.007805/2022

DECLARO, sob as penas da Lei, em atendimento ao Edital do **Pregão Eletrônico nº PE 027/23**, promovido por essa DPRJ, na Avenida Marechal Câmara nº 314, Centro – Rio de Janeiro/RJ CEP: 20.020-080, que a empresa **DECISION SERVICOS DE TECNOLOGIA DA INFORMACAO LTDA**, CNPJ nº **03.535.902/0005-44**, com sede à **RUA GILDASIO AMADO, 55 - SALA 1604 - BARRA DA TIJUCA - RIO DE JANEIRO/RJ - CEP: 22.631-020**, por mim representada, atende ao disposto na Lei 7.258/2016, apresentando um efetivo de **9** empregados.

Rio de Janeiro (RJ), 28 de setembro de 2023.



**ALVARO LUIZ SOARES FUZEIRO**  
**GERENTE DE CONTAS**  
alvaro.fuzeiro@decision-tec.com.br  
Tel. (21) 99649-5376  
ID 05.568.662-0 (IFP/RJ)  
CPF 839.674.197-20

03.535.902/0005-44  
DECISION SERVIÇOS DE  
TECNOLOGIA DA INFORMAÇÃO LTDA  
Rua Gildasio Amado 00055 Sala 1604  
Barra Da Tijuca - Cep: 22631-020  
RIO DE JANEIRO RJ

## ANEXO VII DECLARAÇÃO DE ELABORAÇÃO INDEPENDENTE DE PROPOSTA

DEFENSORIA PÚBLICA DO ESTADO DO RIO DE JANEIRO  
EDITAL DE PREGÃO ELETRÔNICO Nº 027/23  
Processo nº E-20/001.007805/2022

**ALVARO LUIZ SOARES FUZEIRO**, identidade nº **05.568.662-0 (IFP/RJ)** e inscrito no CPF nº **839.674.197-20**, como representante devidamente constituído da empresa **DECISION SERVICOS DE TECNOLOGIA DA INFORMACAO LTDA**, CNPJ nº **03.535.902/0005-44**, com sede à **RUA GILDASIO AMADO, 55 - SALA 1604 - BARRA DA TIJUCA - RIO DE JANEIRO/RJ - CEP: 22.631-020**, doravante denominado LICITANTE, para fins do disposto no Edital nº **027/23**, DECLARA, sob as penas da lei, em especial o art. 299 do código Penal Brasileiro, que:

- a) A proposta anexa foi elaborada de maneira independente, e que o conteúdo da proposta anexa não foi, no todo ou em parte, direta ou indiretamente, informado a, discutido com ou recebido de qualquer outro participante potencial ou de fato do presente certame, por qualquer meio ou por qualquer pessoa;
- b) A intenção de apresentar a proposta anexa não foi informada a, discutida com ou recebida de qualquer outro participante potencial ou de fato do presente certame, por qualquer meio ou qualquer pessoa;
- c) Que não tentou, por qualquer meio ou por qualquer pessoa, influir na decisão de qualquer outro participante potencial ou de fato do presente certame, quanto a participar ou não da referida licitação;
- d) Que o conteúdo da proposta anexa não será, no todo ou em parte, direta ou indiretamente, comunicado ou discutido com qualquer outro participante potencial ou de fato, antes da adjudicação do objeto da referida licitação;
- e) Que o conteúdo da proposta anexa não foi no todo ou em parte, direta ou indiretamente, informado a, discutido com ou recebido de DEFENSORIA PÚBLICA DO ESTADO DO RIO DE JANEIRO antes da abertura oficial das propostas e;



4

# DECISION

SERVIÇOS DE TECNOLOGIA DA INFORMAÇÃO LDO.

f) Que está plenamente ciente do teor e da extensão desta declaração e que detém plenos poderes e informações para firmá-la.

Rio de Janeiro (RJ), 28 de setembro de 2023.



**ALVARO LUIZ SOARES FUZEIRO**

**GERENTE DE CONTAS**

alvaro.fuzeiro@decision-tec.com.br

Tel. (21) 99649-5376

ID 05.568.662-0 (IFP/RJ)

CPF 839.674.197-20

03.535.902/0005-44  
DECISION SERVIÇOS DE  
TECNOLOGIA DA INFORMAÇÃO LTDA  
Rua Gildasio Amado 00055 Sala 1604  
Barra Da Tijuca - Cep: 22631-020  
RIO DE JANEIRO RJ

DECISION

**Brasília (Sede)**

Sector Hoteleiro Sul - Quadra 06 - Conjunto 71  
Bloco A - Sala 102 - Asa Sul - Brasília/DF  
Cep. 70.322-915 - Tel. (61) 3045.0050

**Salvador**

Avenida Tancredo Neves, 620 - Salas 2910 e 2911  
29º andar - Torre Empresarial da Ed. Mundo Plaza  
Caminho das Árvores - Salvador/BA - Cep. 41.820-020  
Tel. (71) 3566.7007

**São Paulo**

Rua Arizona, 1.422 - Conjunto 76 - Ed. Platinum  
Building Bemni - Bemni - São Paulo/SP - Cep. 04.967-003  
Tel. (11) 5583.0344

5

## ANEXO VIII DECLARAÇÃO DE INEXISTÊNCIA DE PENALIDADE

DEFENSORIA PÚBLICA DO ESTADO DO RIO DE JANEIRO  
EDITAL DE PREGÃO ELETRÔNICO Nº 027/23  
Processo nº E-20/001.007805/2022

A empresa **DECISION SERVIÇOS DE TECNOLOGIA DA INFORMAÇÃO LTDA**, inscrita no CNPJ sob o nº **03.535.902/0005-44**, sediada na **RUA GILDASIO AMADO, 55 - SALA 1604 - BARRA DA TIJUCA - RIO DE JANEIRO/RJ - CEP: 22.631-020**, neste ato representada pelo seu representante legal, o Sr. **ALVARO LUIZ SOARES FUZEIRO**, inscrito no CPF sob o nº **839.674.197-20**, portador da cédula de identidade nº **05.568.662-0**, expedida por **IFP/RJ**, DECLARA, sob as penas da Lei, que não foram aplicadas penalidades de suspensão temporária da participação em licitação, impedimento de contratar ou declaração de inidoneidade para licitar e contratar por qualquer Ente ou Entidade da Administração Federal, Estadual, Distrital e Municipal cujos efeitos ainda vigorem.

Rio de Janeiro (RJ), 28 de setembro de 2023.



**ALVARO LUIZ SOARES FUZEIRO**  
GERENTE DE CONTAS  
alvaro.fuzeiro@decision-tec.com.br  
Tel. (21) 99649-5376  
ID 05.568.662-0 (IFP/RJ)  
CPF 839.674.197-20

03.535.902/0005-44  
DECISION SERVIÇOS DE  
TECNOLOGIA DA INFORMAÇÃO LTDA  
Rua Gildasio Amado 00055 Sala 1604  
Barra Da Tijuca - Cep: 22631-020  
RIO DE JANEIRO RJ