

DECISION

SERVIÇOS DE TECNOLOGIA DA INFORMAÇÃO S/A.

Item 2 - PPDM

DECISION

Brasília (Sede)

Setor Hoteleiro Sul - Quadra 06 - Conjunto W
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. 41.820-000
Tel. (71) 3565.7007

São Paulo


Rua Arizona, 1.422 - Conjunto 75 - Ed. Platinum
Building Berrini - Berrini - São Paulo/SP - Cep. 04.367-003
Tel. (11) 5583.0344

PowerProtect Data Manager 19.9

Administration and User Guide

Version 19.9

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.....	11
Chapter 1: Getting Started.....	15
Introducing the PowerProtect Data Manager software.....	15
Supported Internet Protocol versions.....	16
References.....	16
Terminology.....	17
Access the PowerProtect Data Manager UI.....	18
Getting Started window.....	18
UI tools and options	19
Provide customer feedback.....	21
Role-based security.....	22
Chapter 2: Managing Users.....	23
Managing identity providers.....	23
Configure an external identity provider.....	23
Edit an external identity provider.....	24
Delete an external identity provider.....	25
Managing user roles and privileges	25
Managing local identity provider users.....	25
Add a local user.....	25
Edit or delete a local user.....	26
Common password policy.....	26
Change a local user password.....	27
Reset a forgotten local user password.....	27
Reset operating system passwords.....	27
Default authorizations.....	28
System-provided roles and associated privileges.....	28
Role privilege definitions.....	31
External authorization associations.....	34
Add identity provider group-to-role mapping.....	34
Modify identity provider group-to-role mapping.....	34
Delete identity provider group-to-role mapping.....	35
Remote component authentication.....	35
Add a credential.....	35
View credential usage.....	36
Edit a credential.....	36
Delete credentials.....	36
Credential security.....	37
Chapter 3: Managing Storage.....	38
Protection storage.....	38
High Availability PowerProtect DD support.....	38
Add protection storage.....	39
Edit protection storage.....	40

Storage units.....	40
Storage unit limitations.....	41
Storage unit considerations for PowerProtect DD.....	41
Create a storage unit.....	42
Edit a storage unit.....	43
Change a storage unit password.....	44
View the storage unit password.....	45
Overview of PowerProtect Data Manager Cloud Tier.....	45
Chapter 4: Using the PowerProtect Search Engine.....	46
Introducing the PowerProtect Search Engine.....	46
Set up and manage indexing.....	46
Search Engine node deletion.....	48
Delete operational nodes from a Search cluster.....	48
Redeploy or delete failed nodes from a Search cluster.....	49
Edit the network configuration for a search engine node.....	49
Perform a search.....	50
Virtual machine file level restore from a search.....	50
File level restore to original virtual machine using File Search.....	50
File level restore to alternate virtual machine using File Search.....	51
Troubleshooting Search Engine issues.....	52
Chapter 5: Managing Assets.....	58
About asset sources, assets, and storage.....	58
About vCenter Server asset sources and virtual assets.....	58
About other asset sources.....	58
Prerequisites for discovering asset sources.....	59
Enable an asset source.....	60
Disable an asset source.....	61
Delete an asset source.....	61
Adding a vCenter Server asset source.....	61
Add a VMware vCenter Server.....	62
Creating a dedicated vCenter user account.....	63
VM Direct protection engine overview.....	66
Requirements for an external VM Direct Engine.....	66
Add a VM Direct Engine.....	66
Additional VM Direct actions.....	68
Transparent Snapshot Data Mover protection mechanism.....	70
Adding a Cloud Snapshot Manager tenant.....	72
Add a Cloud Snapshot Manager Tenant.....	72
Chapter 6: Managing Protection Policies.....	73
Protection policies.....	73
Before you create a protection policy.....	74
Supported enhanced VMware topologies for virtual-machine protection.....	76
Add a protection policy for virtual-machine protection.....	77
Managing virtual-machine backups.....	84
Add a Cloud Tier schedule to a protection policy.....	86
Manage Cloud Tier asset copies.....	87

Manual backups of protected assets.....	88
Manual replication of protected assets.....	88
Manual Cloud Tiering of protected assets.....	89
Editing a protection policy.....	89
Modify a policy name and description, objectives, or options.....	89
Changing storage targets and storage units.....	90
Add or remove assets in a protection policy.....	91
View assets assigned to a protection policy.....	92
Extended retention.....	93
Edit the retention period for backup copies.....	94
Delete backup copies.....	95
Retry a failed backup copy deletion.....	96
Export data for deleted backup copies.....	96
Remove backup copies from the PowerProtect Data Manager database.....	97
Removing expired backup copies.....	97
Removing assets from PowerProtect Data Manager.....	98
Remove assets and associated protection copies.....	98
Export protection.....	99
Disable a protection policy.....	99
Protection jobs running for a disabled policy.....	100
Enable a disabled protection policy.....	101
Customize the default behavior of disabled policies.....	101
Delete a protection policy.....	101
Add a Service Level Agreement.....	102
Export Asset Compliance.....	104
Protection rules.....	105
Creating virtual machine tags in the vSphere Client.....	105
Add a protection rule.....	106
Manually run a protection rule.....	107
Edit or delete a protection rule.....	108
View assets applied to a protection rule.....	108
Change the priority of an existing protection rule.....	108
Configure protection rule behavior.....	109
Chapter 7: Restoring Data and Assets.....	110
View backup copies available for restore.....	110
Restoring a virtual machine or VMDK.....	111
Restoring a virtual machine backup with the storage policy association.....	111
Prerequisites to restore a virtual machine.....	112
Restore to the original virtual machine.....	112
Restore individual virtual disks.....	114
Restore to a new virtual machine.....	115
Instant access virtual machine restore.....	117
File level restore to original virtual machine.....	120
File level restore to alternate virtual machine.....	121
Direct restore to ESXi.....	122
Restore an application-aware virtual machine backup.....	123
Restore the PowerProtect Data Manager server.....	123
Restore Cloud Tier backups to protection storage.....	124
Recall and restore from Cloud Tier.....	124

Chapter 8: Preparing for and Recovering From a Disaster.....	126
Managing system backups for server disaster recovery.....	126
Server DR protection storage types.....	126
Overview of PowerProtect Data Manager Cloud Disaster Recovery.....	127
Prepare the DD system recovery target (NFS).....	127
Configure PowerProtect Data Manager server DR backups.....	128
Record settings for server DR.....	129
Manage PowerProtect Data Manager server DR backups.....	129
Restore PowerProtect Data Manager from server DR backups.....	130
Recovering the Search Engine from a DR backup.....	131
Troubleshooting NFS backup configuration issues.....	132
Troubleshoot recovery of PowerProtect Data Manager.....	133
Quick recovery.....	133
Quick recovery prerequisites.....	136
Add a remote system for quick recovery.....	137
Edit a remote system.....	137
Quick recovery remote view.....	138
Recover a failed PowerProtect Data Manager backup.....	138
Chapter 9: Managing Alerts, Jobs, and Tasks.....	139
Configure Alert Notifications.....	139
View and manage alerts.....	139
View and manage Audit Logs.....	140
Monitoring jobs and tasks.....	140
Monitor and view jobs.....	141
View details for protection jobs.....	142
View details for system jobs and tasks.....	144
Filter, group, and sort jobs.....	146
Restart a job or task manually.....	148
Restart a job or task automatically.....	148
Resume misfire jobs after a PowerProtect Data Manager update.....	149
Cancel a job or task.....	150
Exporting logs.....	151
Export logs for jobs.....	152
Export logs for assets or tasks.....	152
Chapter 10: Modifying the System Settings.....	153
System settings.....	153
Modify the network settings.....	153
Synchronize time on PowerProtect Data Manager and other systems.....	153
Modify the appliance time zone.....	154
Enable replication encryption.....	154
Backup and restore encryption.....	154
PowerProtect Data Manager licensing.....	156
Specify a vCenter Server as the PowerProtect Data Manager host.....	157
System Support.....	158
Configuring SupportAssist for PowerProtect Data Manager.....	158
Telemetry Collector.....	162

CloudIQ reporting.....	163
Set up the email server.....	163
Add AutoSupport.....	163
Enabling automatic update package checks and downloads.....	164
Add a log bundle.....	164
Audit logging and monitoring system activity.....	164
Monitor system state and system health.....	166
Access the open source software package information.....	166
Security certificates.....	166
Modifying the PowerProtect Data Manager virtual machine disk settings.....	167
Modify the data disk size.....	167
Modify the system disk size.....	168
Memory optimization.....	168
Adjust the memory.....	169
Configure the DD system.....	169
Virtual networks (VLANs).....	170
Supported scenarios.....	171
Virtual network prerequisites.....	171
Configuring virtual networks.....	172
Virtual network asset assignment.....	174
Chapter 11: Protecting Virtual Machines using the Transparent Snapshot Data Mover	177
Overview of transparent snapshots for virtual machine protection.....	177
VIB installation monitoring and management.....	177
Transparent snapshot data mover system requirements.....	178
Prerequisites to virtual machine protection with the Transparent Snapshot Data Mover.....	178
Additional privileges required for a dedicated vCenter user account to use Transparent Snapshot Data Mover.....	178
Creating VMkernel ports.....	179
Virtual machine transparent snapshot unsupported features and limitations.....	180
Transparent Snapshot Performance and Scalability.....	181
Chapter 12: PowerProtect Functionality Within the vSphere Client.....	182
PowerProtect functionality within the vSphere Client.....	182
Overview of the PowerProtect plug-in for the vSphere Client.....	182
Prerequisites for enabling the vSphere Client PowerProtect plug-in.....	183
Monitor PowerProtect Data Manager virtual machine protection copies.....	184
Manual PowerProtect policy backup in the vSphere Client.....	185
Image-level restore of a PowerProtect backup in the vSphere Client.....	185
File-level restore of a PowerProtect backup in the vSphere Client.....	186
Overview of VASA and VMware Storage Policy Based Management	188
Register the VASA provider for policy association.....	188
Add an SPBM policy and associate with a PowerProtect Data Manager virtual machine policy.....	189
Monitor virtual machine protection policy compliance.....	190
Chapter 13: VMware Cloud (VMC) on Amazon Web Services (AWS).....	191
PowerProtect Data Manager image backup and recovery.....	191
Supported PowerProtect Data Manager and DDVE deployment configurations.....	191
Deployment and configuration best practices and requirements.....	192

Configuring the VMC-on-AWS portal.....	192
Interoperability with PowerProtect Data Manager features.....	193
vCenter server inventory requirements.....	193
Creating a dedicated cloud-based vCenter user account.....	193
Specify the required privileges for a dedicated cloud-based vCenter user account	193
Add a VM Direct Engine.....	195
Unsupported operations	197
Chapter 14: Azure VMware Solution (AVS) on Microsoft Azure.....	198
PowerProtect Data Manager image backup and recovery.....	198
Supported PowerProtect Data Manager and DDVE deployment configurations.....	198
Deployment and configuration best practices and requirements.....	199
Configuring the AVS-on-Azure portal.....	199
vCenter server inventory requirements.....	200
Creating a dedicated cloud-based vCenter user account.....	200
Specify the required privileges for a dedicated cloud-based vCenter user account	200
Add a VM Direct Engine.....	202
Unsupported operations	203
Chapter 15: Google Cloud VMware Engine (GCVE) on Google Cloud Product (GCP).....	204
PowerProtect Data Manager image backup and recovery.....	204
Supported PowerProtect Data Manager and DDVE deployment configurations.....	204
Deployment and configuration best practices and requirements.....	205
Configuring the GCVE-on-GCP portal.....	205
vCenter server inventory requirements.....	206
Creating a dedicated cloud-based vCenter user account.....	206
Specify the required privileges for a dedicated cloud-based vCenter user account	206
Add a VM Direct Engine.....	208
Unsupported operations.....	209
Chapter 16: Performing Updates.....	210
Managing update packages.....	210
Automatically check for an update package.....	210
Troubleshooting automatic downloads.....	211
Manually check for an update package.....	211
Download an update package.....	211
Upload an update package.....	212
Delete an update package.....	212
Perform a precheck on an update package.....	212
Install an update package.....	213
Updating the version of PowerProtect Data Manager.....	213
Update PowerProtect Data Manager from version 19.8 to version 19.9.....	214
Update PowerProtect Data Manager from version 19.7 to version 19.9.....	216
Update PowerProtect Data Manager from versions 19.3–19.6 to version 19.9.....	218
Run a manual precheck.....	221
Lockbox passphrase required when updating from some versions.....	221
Chapter 17: Configuring and Managing the PowerProtect Agent Service	223
About the PowerProtect agent service.....	223

Start, stop, or obtain the status of the PowerProtect agent service.....	224
Register the PowerProtect agent service to a different server address.....	224
Recovering the PowerProtect agent service from a disaster.....	225
Restore the PowerProtect Data Manager agent service datastore.....	225
Chapter 18: Backing Up and Recovering a vCenter Server.....	227
Backing up and recovering a vCenter server.....	227
vCenter deployments overview.....	227
Protecting an embedded PSC.....	227
Direct restore to ESXi.....	228
Protecting external deployment models.....	229
vCenter server appliance(s) with one external PSC where PSC fails.....	229
vCenter server appliance is lost but the PSC remains.....	230
vCenter server appliance with multiple PSCs where one PSC is lost, one remains.....	230
vCenter server appliance remains but all PSCs fail.....	230
vCenter server appliance remains but multiple PSCs fail.....	230
vCenter server appliance fails.....	231
vCenter server restore workflow.....	232
Platform Services Controller restore workflow.....	233
Additional considerations.....	233
Command reference.....	234
Chapter 19: Backing Up VMware Cloud Foundation (VCF) on VxRail.....	235
Backing up VCF on VxRail.....	235
VCF and VxRail overview.....	235
VCF components and backup methods.....	236
Check VMware certification.....	237
Backup prerequisites.....	237
The backup script.....	237
Quick protection.....	237
Selective protection: SDDC and NSX-T Managers.....	239
Selective protection: vCenter servers.....	240
Selective protection: vRSLCM, VxRail Manager, Workspace ONE Access, and vRealize Suite virtual machines.....	241
SFTP password change: SDDC and NSX-T Managers.....	241
SFTP password change: vCenter servers.....	242
Backup-script troubleshooting.....	243
Chapter 20: Best Practices and Troubleshooting.....	245
Base 10 standard used for size calculations in the PowerProtect Data Manager UI.....	245
Best practices and additional considerations for the VM Direct Engine.....	245
Change the limit of instant access sessions.....	245
Configuring a backup to support vSAN datastores.....	246
Configuration checklist for common issues.....	246
Disable vCenter SSL certificate validation.....	246
File-level restore and SQL restore limitations.....	247
FLR Agent for virtual machine file level restore.....	248
FLR-supported platform and OS versions for virtual machine restores.....	250
PowerProtect Data Manager resource requirements in a VMware environment.....	251

Software and hardware requirements.....	251
Support for backup and restore of encrypted virtual machines.....	252
Transport mode considerations.....	252
Virtual disk types supported.....	253
Virtual machine data change rate.....	253
VM Direct Engine data ingestion rate.....	254
VM Direct Engine limitations and unsupported features.....	254
VM Direct Engine performance and scalability.....	257
VM Direct Engine selection with virtual networks (VLANs).....	258
Best practices for vCenter Server backup and restore.....	258
Changing the vCenter server FQDN.....	259
Change the vCenter server FQDN.....	258
Monitoring storage capacity thresholds.....	259
Replacing security certificates.....	260
Replacing the self-signed security certificates.....	260
Replace expired or changed certificates on an external server.....	260
Restarting PowerProtect Data Manager.....	262
Scalability limits for vCenter Server, VM Direct Engine and DD systems.....	262
Troubleshooting network setup issues.....	263
Troubleshooting PowerProtect agent service installations.....	263
Troubleshooting PowerProtect agent service operations.....	263
Troubleshoot the PowerProtect agent service operations.....	263
Troubleshooting PowerProtect Data Manager software updates.....	264
Managing certificates after updating PowerProtect Data Manager from versions earlier than 19.1.....	264
Troubleshooting storage units.....	265
Troubleshooting virtual machine backup issues.....	265
Backup completes with a non-quieted snapshot warning.....	265
Backup fails when names include special characters.....	266
Deleting vCenter asset sources or moving ESXi to another vCenter.....	266
Failed to lock Virtual Machine for backup: Another EMC vProxy operation 'Backup' is active on VM.....	267
Lock placed on virtual machine during backup and recovery operations continues for 24 hours if VM Direct appliance fails.....	267
Managing command execution for VM Proxy Agent operations on Linux.....	268
PowerProtect plug-in and portlet for vSphere display errors after replacing security certificates.....	268
SQL databases skipped during virtual machine transaction log backup.....	268
SQL Server application-aware backup displays an error about disk.EnableUUID variable.....	269
SQL Server application-consistent backups fail with error "Unable to find VSS metadata files in directory".....	269
Trailing spaces not supported in SQL database names.....	269
VMware knowledge base articles and product documentation.....	269
Troubleshooting virtual machine restore issues.....	269
Troubleshooting instant access restore failures.....	271
VMware knowledge base articles and product documentation.....	272
Troubleshooting vSphere Plugin deployments.....	272
Troubleshoot vSphere Plugin deployments.....	272
VMware knowledge base articles and product documentation.....	272

Preface

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.

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. In many cases the user interface has not yet been updated to reflect this change.

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.

Purpose

This document describes how to configure and administer Dell EMC 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 Dell EMC PowerProtect Data Manager software.

Revision history

The following table presents the revision history of this document.

Table 1. Revision history

Revision	Date	Description
02	January, 2022	Updated memory requirements and guidance.
01	September, 2021	Initial release of this document for PowerProtect Data Manager version 19.9.

Compatibility information

Software compatibility information for the PowerProtect Data Manager software is provided at the eLab Navigator.

Related documentation

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

- *PowerProtect Data Manager Administration and User Guide*—Describes how to configure the software.
- *PowerProtect Data Manager Deployment Guide*—Describes how to deploy the software.
- *PowerProtect Data Manager Licensing Guide*—Describes how to license the software.

- *PowerProtect Data Manager Release Notes*—Contains information on new features, known limitations, environment, and system requirements for the software.
- *PowerProtect Data Manager Security Configuration Guide*—Contains security information.
- *PowerProtect Data Manager AWS Deployment Guide*—Describes how to deploy the software to Amazon Web Services (AWS).
- *PowerProtect Data Manager Azure Deployment Guide*—Describes how to deploy the software to Microsoft Azure.
- *PowerProtect Data Manager GCP Deployment Guide*—Describes how to deploy the software to Google Cloud Platform (GCP).
- *PowerProtect Data Manager Cloud Disaster Recovery Administration and User Guide*—Describes how to deploy Cloud DR, protect VMs in the AWS or Azure cloud, and run recovery operations.
- *PowerProtect Data Manager for Cyber Recovery User Guide*—Describes how to install, update, patch, and uninstall the Dell EMC PowerProtect Cyber Recovery software.
- *PowerProtect Data Manager for File System Agent User Guide*—Describes how to configure and use the software with the File System agent for file system data protection.
- *PowerProtect Data Manager for Kubernetes User Guide*—Describes how to configure and use the software to protect and recover namespaces and PVCs in a Kubernetes cluster.
- *PowerProtect Data Manager for Microsoft Application Agent Exchange Server User Guide*—Describes how to configure and use the software to protect and recover the data in a Microsoft Exchange Server environment.
- *PowerProtect Data Manager for Microsoft Application Agent SQL Server User Guide*—Describes how to configure and use the software to protect and recover the data in a Microsoft SQL Server environment.
- *PowerProtect Data Manager for Oracle RMAN Agent User Guide*—Describes how to configure and use the software to protect and recover the data in an Oracle Server environment.
- *PowerProtect Data Manager for SAP HANA Agent User Guide*—Describes how to configure and use the software to protect and recover the data in an SAP HANA Server environment.
- *PowerProtect Data Manager for Storage Direct Agent User Guide*—Describes how to configure and use the software with the Storage Direct agent to protect data on VMAX storage arrays through snapshot backup technology.
- *PowerProtect Data Manager for Network Attached Storage User Guide*—Describes how to configure and use the software to protect and recover the data on network attached storage (NAS) shares and appliances.
- *PowerProtect Data Manager Public REST API documentation*—Contains the PowerProtect Data Manager APIs and includes tutorials to guide you in their use.

Typographical conventions

The following type style conventions are used in this document:

Table 2. 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

- The Customer Support website
- The Community Network

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.

Knowledgebase

The Knowledgebase contains applicable solutions that you can search for either by solution number (for example, KB000xxxxxx) or by keyword.

To search the Knowledgebase:

1. Go to the Customer Support website.
2. On the **Support** tab, click **Knowledge Base**.
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 DPAD.Doc.Feedback@emc.com.

Getting Started

Topics:

- Introducing the PowerProtect Data Manager software
- Supported Internet Protocol versions
- References
- Terminology
- Access the PowerProtect Data Manager UI
- Provide customer feedback
- Role-based security

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:

- 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 Dell EMC Data Protection portfolio to enable data protection as a service, providing the following benefits:

- Enables the data protection team to create data paths with provisioning, automation, and scheduling to embed protection engines into the 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:
 - VMware virtual machines
 - File systems
 - VMAX storage groups
 - Kubernetes clusters
 - Microsoft Exchange and SQL databases
 - Oracle databases
 - SAP HANA databases
 - Network attached storage (NAS) shares
- Uses an agent-based approach to discover the protected and unprotected databases on an application server.
- Enables governed self-service and centralized protection by:
 - Monitoring Service Level Objectives (SLOs)
 - Identifying violations of Recovery Point Objectives (RPOs)
- Supports deploying an external VM Direct appliance to move data with a VM Direct Engine. The PowerProtect Data Manager software comes pre-bundled with an embedded VM Direct Engine, which is automatically used as a fallback proxy for performing backup and restore operations when the external VM Direct Engines fail or are disabled. Dell EMC recommends that you deploy external VM Direct Engines, because the embedded VM Direct Engine has limited capacity for performing backup streams. The embedded VM Direct Engine is sufficient, however, for virtual machine crash-consistent protection policies that use the Transparent Snapshot Data Mover (TSDM) protection mechanism.
- Supports the vRealize Automation DP extension, which enables provisioning of virtual machines with PowerProtect Data Manager protection, on-demand backup, and restore to the original or a new location. The *vRealize Automation Data Protection Extension for PowerProtect Data Manager Installation and Administration Guide* provides more information.
- Supports integration of Dell EMC Cloud Disaster Recovery (Cloud DR), including workflows for Cloud DR deployment, protection, and recovery operations in the AWS or Azure cloud.

- Supports PowerProtect Search, which enables backup administrators to quickly search for and restore VM file copies. The Search Service can be enabled by adding a search node to the configurable Search Engine that is autodeployed during the PowerProtect Data Manager deployment.
- Provides a RESTful interface that allows the user to monitor, configure, and orchestrate PowerProtect Data Manager. Customers can use the APIs to integrate their own automation framework or quickly write new scripts with the help of easy-to-follow tutorials.
- Integrates with Dell EMC PowerProtect Cloud Snapshot Manager to view PowerProtect Cloud Snapshot Manager jobs, alerts, and reports from a consolidated PowerProtect Data Manager dashboard.

Supported Internet Protocol versions

PowerProtect Data Manager only supports the use of IPv4 addresses.

Using an IPv6 address can result in errors or other unexpected behavior. When configuring devices to connect over the network with PowerProtect Data Manager, use only IPv4 addresses.

References

Some procedures in this document reference other publications for further details. Additionally, updates to documentation after initial publication are provided in the release notes.

The following publications, available on Customer Support, provide additional product information:

- *PowerProtect Data Manager Administration and User Guide*—Describes how to configure the software.
- *PowerProtect Data Manager Deployment Guide*—Describes how to deploy the software.
- *PowerProtect Data Manager Licensing Guide*—Describes how to license the software.
- *PowerProtect Data Manager Release Notes*—Contains information on new features, known limitations, environment, and system requirements for the software.
- *PowerProtect Data Manager Security Configuration Guide*—Contains security information.
- *PowerProtect Data Manager AWS Deployment Guide*—Describes how to deploy the software to Amazon Web Services (AWS).
- *PowerProtect Data Manager Azure Deployment Guide*—Describes how to deploy the software to Microsoft Azure.
- *PowerProtect Data Manager GCP Deployment Guide*—Describes how to deploy the software to Google Cloud Platform (GCP).
- *PowerProtect Data Manager Cloud Disaster Recovery Administration and User Guide*—Describes how to deploy Cloud DR, protect VMs in the AWS or Azure cloud, and run recovery operations.
- *PowerProtect Data Manager for Cyber Recovery User Guide*—Describes how to install, update, patch, and uninstall the Dell EMC PowerProtect Cyber Recovery software.
- *PowerProtect Data Manager for File System Agent User Guide*—Describes how to configure and use the software with the File System agent for file system data protection.
- *PowerProtect Data Manager for Kubernetes User Guide*—Describes how to configure and use the software to protect and recover namespaces and PVCs in a Kubernetes cluster.
- *PowerProtect Data Manager for Microsoft Application Agent Exchange Server User Guide*—Describes how to configure and use the software to protect and recover the data in a Microsoft Exchange Server environment.
- *PowerProtect Data Manager for Microsoft Application Agent SQL Server User Guide*—Describes how to configure and use the software to protect and recover the data in a Microsoft SQL Server environment.
- *PowerProtect Data Manager for Oracle RMAN Agent User Guide*—Describes how to configure and use the software to protect and recover the data in an Oracle Server environment.
- *PowerProtect Data Manager for SAP HANA Agent User Guide*—Describes how to configure and use the software to protect and recover the data in an SAP HANA Server environment.
- *PowerProtect Data Manager for Storage Direct Agent User Guide*—Describes how to configure and use the software with the Storage Direct agent to protect data on VMAX storage arrays through snapshot backup technology.
- *PowerProtect Data Manager for Network Attached Storage User Guide*—Describes how to configure and use the software to protect and recover the data on network attached storage (NAS) shares and appliances.
- *PowerProtect Data Manager Public REST API documentation*—Contains the PowerProtect Data Manager APIs and includes tutorials to guide you in their use.

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 3. 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 (DDBEA) for databases and applications.
Application Aware	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 SQL databases. You can also schedule 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 VMs, 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 Dell EMC secure Elastic Cloud Storage appliances.
Copy	A PowerProtect Data Manager copy is a point-in-time backup copy of an Asset.
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 VM backup copies can be accessed, mounted, and booted directly from the Protection Storage targets as running VMs. Copies can also be moved to a production VMware datastore using vMotion. PowerProtect Data Manager VM application-aware backup copies can be mounted directly from the Protection Storage targets as running SQL databases, which includes the ability to roll forward log backups. These SQL 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 type, assets, backup start/stop time, backup device, 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 that your organization requires. SLAs are made up of one or more Service Level Objectives.
Service Level Objectives (SLOs)	Definable rules that set the criteria for Recovery Point Objectives (RPOs), encryption, and locations of backups according to your 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.

Steps

1. From a host that has network access to the virtual appliance, use 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.

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 **Getting Started** page appears.

- 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.

Getting Started window

The **Getting Started** window provides configuration options that are required when the system is first deployed.

This window appears upon first deployment of PowerProtect Data Manager and opens to this page by default until you click **Skip This**.

You can access the **Getting Started** page at any time by clicking  and then selecting **Getting Started**.

Table 4. PowerProtect Data Manager Getting Started menu items

Options	Description
Support	View and configure SupportAssist, Email Setup, AutoSupport, Logs, System Health.
Disaster Recovery Backup	Configure and manage backups for disaster recovery.
VMware vCenter	Opens the Infrastructure > Asset Sources window, where you can add a vCenter instance as an asset source so that virtual machine assets can be added to a protection policy.
Protect Assets	Opens the Protection > Protection Policies window, where you can manage protection policy workflows for all asset types.

UI tools and options





Learn about the available tools in the UI.

PowerProtect Data Manager UI tools

Table 5. PowerProtect Data Manager tools

Menu item	Description
 Dashboard	<p>Provides a high-level view of the overall state the PowerProtect Data Manager system and includes the following information:</p> <ul style="list-style-type: none">• Alerts—System alerts• Protection—Details about protection policies• Jobs—Details and status of system and protection jobs. You can use the Protection Jobs and System Jobs windows to manage jobs, search, and view details. Filter and sort the information that appears to find specific jobs or tasks.• Policy—Details include number of successes, failures, and excluded assets for each asset type• Protection Storage—Protection storage usage statistics• Restore—Restore statistics• Compliance—Compliance verification statistics. By default, the in compliance asset count and out of compliance asset count displays for all asset types. You can select a specific asset type from the Asset Type list to display compliance statistics for only that category. PowerProtect Data Manager refreshes the data hourly unless you run an ad hoc discovery.
 Infrastructure	<p>Click Infrastructure to:</p> <ul style="list-style-type: none">• View and manage all assets:<ul style="list-style-type: none">◦ VMware virtual machines◦ File systems◦ VMAX storage Groups◦ Kubernetes clusters◦ Microsoft Exchange◦ SQL databases◦ Oracle databases◦ SAP HANA databases• Add 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 Dell EMC Cloud Disaster Recovery.• Create and manage a Search Cluster.• Add PowerProtect Cloud Snapshot Manager tenants as asset sources for jobs, alerts, and reports.
 Protection	<p>Click Protection to:</p> <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.
 Restore	<p>Click Restore to:</p> <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 5. PowerProtect Data Manager tools (continued)

Menu item	Description
 Alerts	Click Alerts to: <ul style="list-style-type: none"> • View and acknowledge alerts and events. • View and examine Audit logs. • Export audit logs to CSV files. • Set audit log boundaries.
 Administration	Click Administration to: <ul style="list-style-type: none"> • Configure users and roles. • Set password credentials and manage key chains. • View certificates. • Configure alert notifications. • Add LDAP Identity Sources.
 Jobs	Click Jobs to manage jobs, view by protection or system, filter, and view details.
 Reporting	Click Reporting to log in to CloudIQ.

Banner UI options

The following table describes the icons that are located in the PowerProtect Data Manager banner.


Table 6. Banner UI options

Option	Description
	Click to enter search criteria to find assets, jobs, logs, and alerts.
	Click to see recent alerts.
	Click to restore assets from replicated copies through quick recovery. This icon only appears when this system receives replicated metadata from a source system.
	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 Getting Started page.
	Click to log out, and log in as a different user.
	Click to see PowerProtect Data Manager version information.
	Click to obtain more information about PowerProtect Data Manager, access Customer Support, send feedback, or view the REST API documentation.
	Click to launch Cloud Snapshot Manager.

Provide 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 with PowerProtect Data Manager.

Steps

1. Log in to the PowerProtect Data Manager UI.
2. From the banner, click , and then select **Send Feedback**.

The customer feedback survey opens in a new window, as shown in the following figure:



The screenshot shows a survey form with the following elements:

- Dell Technologies** logo at the top left.
- Question: "How satisfied are you with PowerProtect Data Manager?"
- A horizontal scale from 0 to 10. Below the scale, "Very dissatisfied" is written under 0 and "Very satisfied" is written under 10.
- Question: "What would you recommend to make your PowerProtect Data Manager experience better?"
- A large empty text box for the recommendation.
- Text: "In certain cases, Dell may want to follow up with you regarding your comment. If you are willing to participate, please provide your email address below. (Note: this is optional, and your email will not be used for marketing purposes)"
- Label: "Email:"
- An empty text box for the email address.
- A button labeled "SEND MY COMMENT" centered below the email box.
- At the bottom left, a link for "Privacy Policy".
- At the bottom right, the "Powered by iperceptions" logo.

Figure 1. Customer feedback survey

3. (Optional) Complete the fields in the customer feedback survey, and when finished, click **Send My Comment**. 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 Dell can follow up with you regarding your feedback.

 **NOTE:** Customer contact information will not be used for marketing purposes.

Role-based security

PowerProtect Data Manager provides predefined user roles that control access to areas of the user interface and to protected operations. Some of the functionality in this guide 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.

Managing Users

Topics:

- Managing identity providers
- Managing user roles and privileges
- Managing local identity provider users
- Default authorizations
- System-provided roles and associated privileges
- Role privilege definitions
- External authorization associations
- Remote component authentication
- Credential security

Managing identity providers

You can configure an external identity provider that manages usernames and passwords.

Only the Administrator and the Security Administrator roles can manage external identity providers. Manage identity providers and roles through the **Administration > Access Control** pane.

Configure an external identity provider

Only the Administrator and the Security Administrator roles can configure an external identity provider.

Steps

1. From the left navigation pane, select **Administration > Access Control**.
The **Access Control** window appears.
2. Click the **Directory Settings** tab.
PowerProtect Data Manager displays a list of configured identity providers.
3. Click **Add**.
The **Add Directory** window appears.
4. Configure the following attributes:

Table 7. Identity provider attributes

Attribute	Description
Server Type	Select a supported identity provider type.
Server Address	Type the hostname or IP address of the identity provider. A protocol prefix is not required.
Secure Connection	Select this attribute if the identity provider uses a secure connection method such as LDAPS or AD over SSL. Selecting this attribute enables the certificate validation controls.
Port	Type the port number for the identity provider.
Domain	Type the domain for which this identity provider authenticates users. For example, <code>ldap.example.com</code> .
User Name	Type a user account that has full read access to the directory. A domain is not required.
Password	Type the password for the specified user account.

Table 7. Identity provider attributes (continued)

Attribute	Description
Group Search Attribute	Type the attribute name that the identity provider should use to validate the group name in the hierarchy.
Group Member Attribute	Type the attribute name that the identity provider should use to validate the group member in the hierarchy.
Group Search Base	If searches should not start from the default base, type the name of a base from which searches should start. Otherwise, leave this attribute empty. Separate multiple search bases with semicolons.

Populate the default values from this table into the appropriate fields when indicated:

Table 8. Default attribute values

Attribute	Value or format	
	AD and AD over SSL	LDAP and LDAPS
Port	<ul style="list-style-type: none"> For unsecure connections, the default port number is 389. For secure connections, the default port number is 636. 	
Group Search Attribute	sAMAccountName	cn
Group Member Attribute	member	memberUid

5. If you selected a secure connection method:
 - a. Click **Verify**.
 - b. In the **Verify Certificate** window, verify the details of the identity provider TLS certificate and then click **Accept**.

NOTE: When you specify the LDAPS protocol, PowerProtect Data Manager automatically downloads the certificates required to connect to the identity provider. Once downloaded, the **Certificate Validation** field appears. Click **Verify** to compare the displayed certificate information with the expected certificate information. If the certificates match, click **Accept** to continue with the setup. Otherwise, click **Cancel** to cancel the setup.
6. Click **Save**.


Next steps

Assign identity provider groups to a role. The section *Add identity provider group-to-role mapping* on page 34 provides instructions. You cannot log in as an external user without mapping users or groups to roles.

Edit an external identity provider

Only the Administrator and the Security Administrator roles can edit an external identity provider.

Steps

1. From the left navigation pane, select **Administration > Access Control**.
The **Access Control** window appears.
2. Click the **Directory Settings** tab.
PowerProtect Data Manager displays a list of configured identity providers.
3. To view more information about an identity provider, click  in the **Details** column for that identity provider.
PowerProtect Data Manager opens the **Details** pane, which displays information about the identity provider's configuration.
4. Select the identity provider, and then click **Edit**.
5. Edit the attributes as required.
6. Click **Save**.

Delete an external identity provider

Only the Administrator and the Security Administrator roles can delete an external identity provider.

Steps

1. From the left navigation pane, select **Administration > Access Control**.
The **Access Control** window appears.
2. Click the **Directory Settings** tab.
PowerProtect Data Manager displays a list of configured identity providers.
3. Select the identity provider that you would like to delete, and then click **Delete**.

Managing user roles and privileges

Users are defined by the local identity provider or by an external identity provider. Users and groups can access all protection policies and assets within the PowerProtect Data Manager environment.

A user's assigned role defines the associated privileges and determines the tasks that the user can perform.

Managing local identity provider users

The following topics describe basic user management for the local identity provider. Only the Administrator and the Security Administrator roles can manage users. The Administrator, Security Administrator, and User roles can view users.

An identity provider is an abstract source of user and group data that PowerProtect Data Manager can map to corresponding roles. An identity provider can be internal to PowerProtect Data Manager or external, such as a supported directory service. PowerProtect Data Manager queries an identity provider to authenticate users as part of the log-in process.

The *PowerProtect Data Manager Security Configuration Guide* provides more information about identity providers, including configuration, role-mapping, and external users. The *PowerProtect Data Manager Security Configuration Guide* also provides information about the local identity provider.

NOTE: User authorization grants or denies users access to PowerProtect Data Manager resources. Authorization is the same for local identity provider users and external identity provider users.

You can create local users to perform management tasks. When you create a local user account, you must assign a role to the user.

Add a local user

Only the Administrator and the Security Administrator roles can add users to the local identity provider.

Steps

1. From the left navigation pane, select **Administration > Access Control**.
The **Access Control** window appears.
2. Click the **Users/Groups** tab.
PowerProtect Data Manager displays a list of configured user accounts and external identity provider groups, including any associated roles.
3. Click **Add User/Group**.
The **Add User/Group** window opens.
4. Select **Local User**.
5. Provide the following information:
 - **First Name**
 - **Last Name**
 - **User Name**
 - **Email Address**
 - **Password**

- Retype to confirm the password.
- **Force Password Change**—Enabled by default. Requires the user to update the password at first login.
- **Role**

8. Click **Save**.

Results

The new user appears in the list of configured user accounts.

Edit or delete a local user

Only the Administrator and the Security Administrator roles can edit or delete local identity provider users.

Steps

1. From the left navigation pane, select **Administration > Access Control**.
The **Access Control** window appears.
2. Click the **Users/Groups** tab.
PowerProtect Data Manager displays a list of configured user accounts and external identity provider groups, including any associated roles.
3. Click  for any user account to see the following information:
 - Username
 - First name
 - Last name
 - Email address
 - User role
 - Date the user was created
4. Select the user that you want to edit or delete.
5. Do one of the following:
 - To delete the user, click **Delete**.
 - To edit the user, click **Edit**, modify the user fields, and then click **Save**.

Results

The changes appear in the list of configured user accounts.

Common password policy

When you set a local identity provider account password, ensure that the credential meets the following requirements:

- Contains a minimum of nine characters and a maximum of one hundred characters
- Contains at least one numeric character (0-9)
- Contains at least one uppercase character (A-Z)
- Contains at least one lowercase character (a-z)
- Contains at least one special character from the following list of acceptable characters:

```
!@#$%^&*()_+~{}|[]<>?/'":;.\`"
```

Spaces are allowed.

- Contains only letters from the English alphabet
- Does not contain other sensitive information that is associated with the user account, such as the first and last names, username, or email address

Change a local user password

Use the self-service feature to change the password for a local identity provider user.

Prerequisites

If you do not know the current password, [Reset a forgotten local user password on page 27](#) provides more information. External identity provider users cannot reset their password using this procedure. Contact the identity provider administrator to reset your password.

Steps

1. Log in to the PowerProtect Data Manager UI.
2. From the banner, select **User Options > Change Password**.
3. Type the current password for the local user.
4. Type the new password twice for confirmation.
The new password must conform to the [Common password policy on page 26](#).
5. Click **Save**.

Reset a forgotten local user password

Use the self-service feature to reset a forgotten password for a local user.

Prerequisites

- The account must be a local identity provider user.
- A mail server must be configured on PowerProtect Data Manager.
- External identity provider users cannot reset their password using this procedure. Contact the identity provider administrator to reset your password.

Review [Common password policy on page 26](#) before you select a new password.

About this task

Local users can receive an email with a link to reset their password. The reset password link in the email expires in 20 minutes, after which time they must request another link.

Steps

1. In the PowerProtect Data Manager login page, click **Forgot Password**.
2. In the **Forgot Password** dialog box, type your user name, click **Send Link**, and click **OK** to dismiss the informational dialog box.
The system sends a message to the email address associated with your user name.
3. Open the email and click the link.
4. In the **Reset Password** dialog box, type a new password in the **New Password** and **Confirm New Password** fields, and click **Save**.
The PowerProtect Data Manager login page appears.
5. Log in with your user name and new password.

Reset operating system passwords


Only the Administrator role can reset operating system passwords. You can change the password for the Linux operating system root, admin, and support users by using the PowerProtect Data Manager UI.

About this task

For the root user, this method works if the current password has not expired and you know the current password. If the password has expired, the attempt fails.

Review [Common password policy on page 26](#) before you select a new password.

Steps

1. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.
2. Click  and then select **Authentication**.
The **System Users** window displays.
3. Select the password you want to change:
 - For the root and support users, click **Edit**.
 - For the operating system admin user, click **Reset**. You can reset the operating system admin user password without providing the existing password.
4. Update the form, and then click **Save**.

Default authorizations

Take note of the following user, group, and role considerations when authorizing users or adding users to roles and groups.

Default admin user

The default admin user is preassigned the Administrator role during PowerProtect Data Manager installation.

The default admin user has super user control over PowerProtect Data Manager and cannot be deleted. However, you can modify the attributes of the default admin user.

Oracle group users

Note that users in the Oracle group have permission to delete the lockbox configuration file. To prevent data loss, add only trusted users to this group.

System-provided roles and associated privileges

A role defines the privileges and permissions that a user has to perform a group of tasks. When a user is assigned a role, you grant the user all of the privileges that are defined by the role.

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

You can assign a user to multiple roles. For example, a user who has both Backup Administrator and Restore Administrator roles but does not have full system administration privileges.

Administrator role

The system Administrator role is responsible for setup, configuration, and all PowerProtect Data Manager management functions. The Administrator role provides systemwide access to all functionality across all organizations. One default Administrator role is assigned at PowerProtect Data Manager deployment and installation. You can add and assign additional Administrator roles to users in your organization who require full access to the system.

User role

The User role is responsible for monitoring the PowerProtect Data Manager Dashboard, Activity Monitor, and Notifications. The User role provides read-only access to monitor activities and operations. Assign the User role to users in your organization who monitor Dashboard activities, Activity Monitor, and Notifications. Users with this role do not require the ability to configure the system or access backup data. Most privileges that are held by this role are read-only.

Security Administrator role

The Security Administrator role is defined for a limited set of users whose manage user accounts and roles, privileges, audit logs, and authentication sources. These functions are separate from the Administrator role. You can assign this role to individuals with security clearances who may not be responsible for day-to-day operations but who clear other users for access.

Backup Administrator role

The Backup Administrator role is responsible for defining, configuring, and completing protection tasks such as backup operations. Individuals with this limited access role do not require the full set of system administrator permissions. These users work with resources that the system administrator has already configured. The Backup Administrator role can backup assets and manage copies at the asset level but cannot back up at the protection policy level.

Restore Administrator role

The Restore Administrator role is responsible for completing restore operations. Individuals with this limited access role do not require the full set of system administrator permissions. These individuals work with backups that exist in protection storage and with resources that the system administrator has already configured.

Role privileges

The following table details the privileges that correspond to each predefined role. Role privilege definitions on page 31 provides more information about the allowed activities for each privilege.

Table 9. Role privileges

Category	Roles				
	Administrator	User	Security Administrator	Backup Administrator	Restore Administrator
Monitoring					
View Events	Y	Y	N	Y	Y
Manage Events	Y	N	N	Y	Y
View Historical Data	Y	Y	N	N	N
View Task/Activities	Y	Y	N	Y	Y
Manage External Notifications	Y	N	N	N	N
Security and System Audit					
View Security/System Audit	Y	Y	Y	N	N
Manage Security/System Audit	Y	N	Y	N	N
User and Security Management					
View User Security	Y	Y	Y	N	N
Manage User Security	Y	N	Y	N	N
Support Assistance and Log Management					
View Diagnostic Logs	Y	Y	N	N	N
Manage Diagnostic Logs	Y	N	N	N	N
System Management					
View System Settings	Y	Y	Y	Y	Y

Table 9. Role privileges (continued)

Category	Roles				
	Administrator	User	Security Administrator	Backup Administrator	Restore Administrator
Manage System Settings	Y	N	N	N	N
Activity Management					
Manage Task	Y	N	N	Y	Y
Workflow Execution	Y	N	N	N	N
Manage Discovery Jobs	Y	N	N	N	N
Asset Management					
View Assets	Y	Y	Y	Y	Y
Manage Assets	Y	N	N	Y	N
View Asset Sources	Y	Y	N	Y	Y
Manage Asset Sources	Y	N	N	N	N
View Host	Y	Y	N	Y	Y
Manage Host	Y	N	N	N	Y
View Protection Engines	Y	Y	N	Y	Y
Manage Protection Engines	Y	N	N	N	N
View Search Engines	Y	Y	N	Y	Y
Manage Search Engines	Y	N	N	N	N
Storage Management					
View Protection Storage Targets	Y	Y	N	Y	Y
Manage Protection Storage Targets	Y	N	N	N	N
View Storage Array	Y	Y	N	Y	Y
Manage Storage Array	Y	N	N	N	N
Manage Network	Y	N	N	N	N
Protection Policy					
View Policies	Y	Y	N	Y	N
Manage Policies	Y	N	N	N	N
Recovery and Reuse Management					
Rollback to Production	Y	N	N	N	Y
Recovery to Alternate Location	Y	N	N	N	Y
Export for Reuse	Y	N	N	N	Y
SLA Compliance Management					
View SLA/SLO	Y	N	N	Y	N
Manage SLA/SLO	Y	N	N	N	N
Copy Management					
View Copies	Y	N	N	Y	Y

Table 9. Role privileges (continued)

Category	Roles				
	Administrator	User	Security Administrator	Backup Administrator	Restore Administrator
Manage Copies	Y	N	N	Y	N
View Retention Range	Y	N	N	Y	N
Manage Retention Range	Y	N	N	N	N
Delete Copies	Y	N	N	N	N
All Copies Search	Y	N	N	N	N
Resource Group					
View Resource Groups	Y	N	Y	N	N
Manage Resource Groups	Y	N	Y	N	N

Role privilege definitions

System-provided roles and associated privileges on page 28 lists the privileges that PowerProtect Data Manager associates with each integrated role. For each privilege, the following tables identify the specific tasks which a user with that privilege can perform.

Table 10. Monitoring privileges

Privilege	Task
View Events	<ul style="list-style-type: none"> View alerts and external notifications.
Manage Events	<ul style="list-style-type: none"> Create, publish, cancel, ignore, promote, and demote alerts and external notifications.
View Historical Data	<ul style="list-style-type: none"> View historical data that relates to plans, arrays, data targets, data sources, and capacity data.
View Tasks or Activities	<ul style="list-style-type: none"> View task resources.
Manage External Notifications	<ul style="list-style-type: none"> Subscribe or unsubscribe a user for alert notifications.

Table 11. Security and system audit privileges

Privilege	Task
View Security/System Audit	<ul style="list-style-type: none"> View security audit-related events and activities.
Manage Security/System Audit	<ul style="list-style-type: none"> Acknowledge security audit-related events and activities. Export audit/change log of events and activities.

Table 12. Support assistance and log management privileges

Privilege	Task
View Diagnostic Logs	<ul style="list-style-type: none"> View log bundle resources. View log information resources. View the log source resource. View logs.
Manage Diagnostic Logs	<ul style="list-style-type: none"> View and manage log bundle resources. View and edit the log source resource. Export logs.

Table 13. User and security management privileges

Privilege	Task
View User Security	<ul style="list-style-type: none"> View users and roles. View identity providers.
Manage User Security	<ul style="list-style-type: none"> Create, view, edit, and delete users. Create, view, edit, and delete roles. Create, view, edit, and delete identity providers. Create, view, edit, and delete user groups.

Table 14. System management privileges

Privilege	Task
View System Settings	<ul style="list-style-type: none"> View server disaster recovery artifacts. View maintenance mode. View license information. View server disaster recovery status. View SupportAssist information. View node, configuration EULA, operating system user, update package, component, configuration status, configuration logs, time zone, and state resources.
Manage System Settings	<ul style="list-style-type: none"> Manage server disaster recovery activities. Manage SupportAssist gateway connection and other telemetry communications. View and edit node state resources. Update license information. View component, configuration status, configuration logs, time zone, and state resources. View and edit node, configuration EULA, operating system user, and lockbox resources. Create, view, edit, and delete update package resources.

Table 15. Activity management privileges

Privilege	Task
Manage Task	<ul style="list-style-type: none"> Create, view, edit, and cancel activity resources.
Workflow Execution	<ul style="list-style-type: none"> Start and cancel workflow execution. View the status of workflow execution.
Manage Discovery Jobs	<ul style="list-style-type: none"> Create, view, edit, and delete discovery jobs.

Table 16. Asset management privileges

Privilege	Task
View Assets	<ul style="list-style-type: none"> View assets.
Manage Assets	<ul style="list-style-type: none"> Create, view, edit, and delete assets.
View Asset Sources	<ul style="list-style-type: none"> View asset sources.
Manage Asset Sources	<ul style="list-style-type: none"> Create, view, edit, and delete asset sources.
View Host	<ul style="list-style-type: none"> View asset hosts.
Manage Host	<ul style="list-style-type: none"> Create, view, edit, and delete asset hosts.
View Protection Engines	<ul style="list-style-type: none"> View protection engines.
Manage Protection Engines	<ul style="list-style-type: none"> Create, view, edit, and delete protection engines.
View Search Engine	<ul style="list-style-type: none"> View the Search Engine.
Manage Search Engine	<ul style="list-style-type: none"> Create, view, edit, and delete the Search Engine.

Table 17. Storage management privileges

Privilege	Task
View Protection Storage Targets	<ul style="list-style-type: none"> View storage targets.
Manage Protection Storage Targets	<ul style="list-style-type: none"> Create, view, edit, and delete storage targets.
View Storage Array	<ul style="list-style-type: none"> View storage arrays.
Manage Storage Array	<ul style="list-style-type: none"> Create, view, edit, and delete storage arrays.
Manage Network	<ul style="list-style-type: none"> Assign network interfaces to storage arrays.

Table 18. Protection policy privileges

Privilege	Task
View Policies	<ul style="list-style-type: none"> View a list of all protection policies. View the storage targets of protection policy. View the accessible assets that are assigned to protection policies. View protection policy schedules. View protection policy networking and other advanced options. View file filters. View protection rules.
Manage Policies	<ul style="list-style-type: none"> Create, view, edit, and delete protection policies. Disable protection policies. Create, view, edit, and delete schedule resources. Add, view, and edit protection policy storage targets. Add, view, and edit protection policy assets. Perform manual backups of protected assets. Create, view, edit, and delete file filters. Create, view, edit, and delete protection rules filters. Assign network interfaces.

Table 19. Recovery and reuse management privileges

Privilege	Task
Rollback to Production	<ul style="list-style-type: none"> Create, view, edit, and start restore to production operations. Create, view, edit, and delete resources that are related to media manager assets.
Recovery to Alternate Location	<ul style="list-style-type: none"> Create, view, edit, and start restore to alternate location operations. Create, view, edit, and delete resources that are related to media manager assets.
Export for Reuse	<ul style="list-style-type: none"> Create, view, edit, and start export and reuse operations. Create, view, edit, and delete resources that are related to media manager assets.

Table 20. SLA compliance management privileges

Privilege	Task
View SLA/SLO	<ul style="list-style-type: none"> View compliance results.
Manage SLA/SLO	<ul style="list-style-type: none"> Create, view, edit, delete, and export compliance results.

Table 21. Copy management privileges

Privilege	Task
View Copies	<ul style="list-style-type: none"> View asset copies and backups.
Manage Copies	<ul style="list-style-type: none"> Edit asset copy and backup retention. Recall copies from the cloud.

Table 21. Copy management privileges (continued)

Privilege	Task
	<ul style="list-style-type: none"> Edit asset copy and backup recall retention.
View Retention Range	<ul style="list-style-type: none"> View retention range.
Manage Retention Range	<ul style="list-style-type: none"> Manage retention range across all copies and backups.
Delete Copies	<ul style="list-style-type: none"> Delete copies and backups.
All Copies Search	<ul style="list-style-type: none"> Manage available copies and backups.

Table 22. Resource group privileges

Privilege	Task
View Resource Groups	<ul style="list-style-type: none"> View a list of all resource groups. View resource group details.
Manage Resource Groups	<ul style="list-style-type: none"> Create, view, edit, and delete resource groups.

External authorization associations

This section describes how to connect PowerProtect Data Manager authorization to identity provider-based subjects.

Add identity provider group-to-role mapping

Only the Administrator and the Security Administrator roles can add identity provider group-to-role mapping.

Steps

- From the left navigation pane, select **Administration > Access Control**.
The **Access Control** window appears.
- Click the **Users/Groups** tab.
PowerProtect Data Manager displays a list of configured user accounts and external identity provider groups, including any associated roles.
- Click **Add User/Group**.
The **Add User/Group** window opens.
- Select **AD/LDAP User Group**.
- Select the domain which corresponds to the identity provider for which you would like to add group-to-role mapping.
- In **Groups**, start typing the name of a identity provider group.
PowerProtect Data Manager searches the identity provider and displays any matching groups.
- Select one or more groups from the list of results.
- Select one or more roles for all group users.
- Click **Add**.


Modify identity provider group-to-role mapping

Only the Administrator and the Security Administrator roles can modify identity provider group-to-role mapping.

Steps

- From the left navigation pane, select **Administration > Access Control**.
The **Access Control** window appears.
- Click the **Users/Groups** tab.

PowerProtect Data Manager displays a list of configured user accounts and external identity provider groups, including any associated roles.

3. Click  for any group to see the following information:
 - Group name
 - Group type
 - Group role
 - Date the group was mapped
4. Select the group that you want to edit, and then click **Edit**.
The **Edit User/Group** window opens.
5. Assign the group to a different role.
The domain and group name are read-only.
6. Click **Save**.

Delete identity provider group-to-role mapping

Only the Administrator and the Security Administrator roles can delete identity provider group-to-role mapping.

Steps

1. From the left navigation pane, select **Administration > Access Control**.
The **Access Control** window appears.
2. Click the **Users/Groups** tab.
PowerProtect Data Manager displays a list of configured user accounts and external identity provider groups, including any associated roles.
3. Select the group that you want to delete, and then click **Delete**.
4. Click **OK** to confirm the deletion.

Remote component authentication

The PowerProtect Data Manager lockbox securely stores known secrets. These secrets include any user account and protection storage credentials that you supply as you configure the software.

Credential security on page 37 provides more information about the lockbox.

PowerProtect Data Manager can use stored credentials in multiple contexts. The term "consumer" means a place where the appliance uses a credential, for any purpose. For example:

- A username and password may apply to one individual host or asset. In this case, the host or asset is the consumer.
- The same credential could also apply to all assets on the same protection policy, if the assets all authenticate with the same username and password. In this case, the protection policy is the consumer, even though the credential applies to the assets under that policy.

You can manage stored credentials through the PowerProtect Data Manager UI or the REST API.

Add a credential

Supply PowerProtect Data Manager with the necessary credentials to access external systems, such as storage targets, assets, and asset sources. You can also add credentials when you create a protection policy.

Steps

1. From the left navigation pane, select **Administration > Credentials**.
The **Credentials** window appears.
2. Click **Add**.
The **Add Credential** dialog box opens.
3. Type a name for the credential.
Credential names should clearly identify the intended purpose and usage.

4. Select a credential type from the drop-down list.
The credential type determines the remaining fields. For example, username and password, token, or key.
5. Complete the remaining fields according to the selected type.
6. Click **Save**.
PowerProtect Data Manager adds the credential to the keystore.

View credential usage

For each stored credential, you can see a list of items that use that credential.

Steps

1. From the left navigation pane, select **Administration > Credentials**.
The **Credentials** window appears.
2. Locate the credential in the list of stored credentials.
Use the filters and column sort options to organize the list of credentials.
3. Select the credential from the list.
Review the **Consumer Count** column for that credential. If the count is zero, the credential is not used anywhere.
4. Select the number in the **Consumer Count** column.
The **Details** pane opens and displays a list of consumers that use the selected credential. The list groups items by type. For example, assets, protection policies, or storage targets.

Edit a credential

You can change a credential name or stored authentication details, such as a username or password. You cannot change the credential type.

Steps

1. From the left navigation pane, select **Administration > Credentials**.
The **Credentials** window appears.
2. Locate the credential in the list of stored credentials.
Use the filters and column sort options to organize the list of credentials.
3. Select the credential from the list, and then click **Edit**.
The **Edit Credential** dialog box opens.
4. Modify any appropriate values.
The available values depend on the credential type. For example, username and password, token, or key.
5. Click **Save**.
PowerProtect Data Manager updates the stored credential.

Delete credentials

You can delete any credentials that are no longer in use or which you no longer need. Deleting a credential creates an entry in the audit log.

Prerequisites

The credentials must not be used anywhere. Verify the credential usage and that the consumer count is zero. If necessary, update anything that uses the credentials, such as protection policies or assets.

Steps

1. From the left navigation pane, select **Administration > Credentials**.
The **Credentials** window appears.
2. Locate the credential in the list of stored credentials.

Use the filters and column sort options to organize the list of credentials.

3. Select the credential or credentials from the list.
4. Verify that the **Consumer Count** column displays zero consumers.
If the count is zero, the credential is not used anywhere and you can delete the credential. The **Delete** button activates when all selected credentials have zero consumers.
5. Click **Delete**.
6. Click **OK** to confirm the deletion.
PowerProtect Data Manager removes the credential.

Credential security

The PowerProtect Data Manager lockbox securely stores known secrets in a central location.

All stored secrets in the lockbox are encrypted. When an activity requires information from the lockbox, the requesting process provides the lockbox passphrase and then receives the required information in a decrypted format.

The lockbox holds secrets such as:

- Credentials for local user accounts.
- Protection storage credentials that you supply as you configure the appliance.
- Credentials by which application agents authenticate to protected assets.

PowerProtect Data Manager creates a strong, unique passphrase during deployment to protect the lockbox contents. After deployment, PowerProtect Data Manager automatically encrypts and manages the lockbox passphrase without user interaction. Automatic management removes the requirement to provide the lockbox passphrase when you update from supported releases. Server DR backups protect the lockbox and its contents.

The File System agent also uses a separate lockbox on protected hosts to store sensitive information, including the credentials by which the application agent accesses external storage infrastructure.

For Kubernetes, PowerProtect Data Manager stores the necessary certificates and credentials for protection operations in a secret resource on the Kubernetes cluster. The Kubernetes documentation provides more information about how to enable encryption for this secret resource.

Managing Storage

Topics:

- Protection storage
- Storage units
- Overview of PowerProtect Data Manager Cloud Tier

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 that manages multiple DD systems

NOTE: Data Domain is now PowerProtect DD. References to Data Domain or DD systems in this documentation, in the UI, and elsewhere in the product include PowerProtect DD systems and older Data Domain systems. In many cases the UI has not yet been updated to reflect this change.

The most up-to-date software compatibility information for PowerProtect Data Manager is provided in the eLab 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.

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

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 on page 39 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 in the eLab Navigator.

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

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**).
4. To add a High Availability PowerProtect DD system, select the checkbox.
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**.
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 external DD systems only in the **Storage** window **Name** column. PowerProtect Data Manager displays PowerProtect DD Management Center storage types in the **Managed By** column.

Edit protection storage

You can change the name, port number, and credentials for an existing protection storage system. You cannot change the address. Only the Administrator role can edit protection storage.

Steps

1. From the left navigation pane, select **Infrastructure > Storage**.
The **Storage** window appears.
2. In the **Protection Storage** tab, select a protection storage system and then click **Edit**.
3. In the **Edit Storage** dialog box, specify the storage system attributes:
 - a. In the **Name** field, specify a new storage name.
 - b. In the **Port** field, specify the port for SSL communication. Default is 3009.
 - c. Under **Host Credentials**, select a new set of credentials or click **Add**.
4. 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**.
5. Click **Save** to exit the **Add Storage** dialog.

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).

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.

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 under the control of 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.

Managing Protection Policies on page 73 provides 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 even if retention lock is enabled.

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 for Oracle RMAN Agent 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 Exchange, the next backup is automatically promoted to a full backup.
 - For SQL, 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 for Oracle RMAN Agent 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 deleted 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.
- The following limitations apply to the number of supported storage units by PowerProtect DD model:

Table 23. Supported storage units for PowerProtect DD Operating System (DDOS) versions

PowerProtect DD system	DDOS version	Maximum number of storage units supported	Supported configurable concurrently active storage units
DD9800	6.0 and later	256	256
DD9600	5.7 and later	256	256
DD6800, DD9300	6.0 and later	128	128

Table 23. Supported storage units for PowerProtect DD Operating System (DDOS) versions (continued)

PowerProtect DD system	DDOS version	Maximum number of storage units supported	Supported configurable concurrently active storage units
DD6300	6.0 and later	100	32
DD990, DD4200, DD4500, DD7200	5.7 and later	128	128
All other DD systems	5.7 and later	100	Up to 32, based on the model
DD9500	5.6	100	64
DD990, DD890	5.3 and later	100	Up to 32, based on the model
DD7200, DD4500, DD4200	5.4 and later	100	Up to 32, based on the model
All other DD systems	5.2 and later	100	Up to 14, based on the model

Table 24. Supported storage units in PowerProtect DD Virtual Edition (DDVE) by TB

Number of TBs	Maximum number of storage units	Supported configurable concurrently active storage units
4	100	6
6		
8		
32	100	14
48		
64	100	32
96		

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

- From the left navigation pane, select **Infrastructure > Storage**.
The **Storage** window appears.
- 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.
- Select **Add**.
The **Create Storage Unit** dialog box opens.
- Type a name for the new storage unit and then select a set of credentials.
Alternatively, you can select **Add Credentials** from the list to add new credentials. Provide a descriptive name for the credentials, a username, and a password. Then, click **Save** to store the credentials.
- 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.

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.
6. Select **Save**.

Results

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



Edit a storage unit

Configure the quota 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** dialog box opens.
5. 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.
 **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.
6. Select **Save**.

Results

PowerProtect Data Manager updates the storage unit quota settings.

Change a storage unit password

It is recommended that you change passwords periodically for security purposes. Change the password for an existing storage unit through the PowerProtect Data Manager UI. The change synchronizes automatically with the protection storage system.

Prerequisites

Before changing a password, verify that recent backup operations have successfully completed by checking the backup status history in the **Jobs** window. You can also perform a new backup of the applicable protection policies.

The default password policy for storage units is:

- Must be between 16 and 20 characters in length
- At least one numeric character (0-9)
- At least one uppercase character (A-Z)
- At least one lowercase character (a-z)
- At least one of the following special characters: \~!@#%&*+={}|:~<>?[]-_,^'/'
- A maximum of three consecutive identical characters

You can modify the password policy for storage units. [Modify the password policy for storage units on page 44](#) provides more information.

Steps

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

The **Storage** window appears.

2. On the **Protection Storage** tab, select a protection storage system, and then select **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 one or more storage units from the list for which you would like to change the password, and then click **Update Password**.

You can only update the password for storage units that are under the control of PowerProtect Data Manager. For other storage units, the **Update Password** button is disabled.

The **Update Password for Storage Unit(s)** dialog box opens.

4. To automatically create a password for the storage unit:
 - a. Select **Automatically generate a new password**.
 - b. Click **Save**.

If you selected multiple storage units, PowerProtect Data Manager creates a unique password for each storage unit.

The storage unit password updates and synchronizes automatically with the protection storage system.

5. To create your own password:
 - a. Select **Enter a new password**.
 - b. In the **Password** field, type a new password for the storage unit according to the password policy.
 - c. Click **Save**.

If you selected multiple storage units, PowerProtect Data Manager uses the same password for each storage unit.

The storage unit password updates and synchronizes automatically with the protection storage system.

6. Go to the **Jobs** window to monitor the progress of the password change operation.
7. Perform another backup of the protection policy and verify that the backup completes successfully.

Modify the password policy for storage units

You can change the rules for passwords that PowerProtect Data Manager applies to storage units.

Steps

1. Connect to the PowerProtect Data Manager console as an admin user.
2. Locate the password policy file at `/usr/local/brs/lib/cbs/config/powerprotect_dd_password_policy.properties`.
3. Modify the password policy.

NOTE: Ensure that the modified password policy complies with the DD password policy.

- Restart the `cbs` service:
`cbs restart`
- Verify that the `cbs` service started successfully.

Results

The modified password policy takes effect when the `cbs` service successfully starts.

View the storage unit password

PowerProtect Data Manager provides a script to retrieve the password for a storage unit that you configured as a backup target.

Prerequisites

This task requires the name of the PowerProtect DD MTree where the storage unit resides.

Steps

- Connect to the PowerProtect Data Manager console as an admin user.
- Navigate to the `/usr/local/brs/puppet/scripts` directory.
- Obtain the storage unit password by typing the following command:
`./get_dd_mtree_credential.py MTree-name`

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.

Using the PowerProtect Search Engine

Topics:

- Introducing the PowerProtect Search Engine
- Set up and manage indexing
- Search Engine node deletion
- Edit the network configuration for a search engine node
- Perform a search
- Virtual machine file level restore from a search
- Troubleshooting Search Engine issues

Introducing the PowerProtect Search Engine

When you install 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 search cluster. Adding a 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. Recovering the Search Engine from a DR backup on page 131 provides instructions. The indexing recovery process will be automated in a future release.

When a DR backup is run, scheduled, or manually triggered, the search 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 ServerDR and Search cluster backups.

Limitations

PowerProtect Search 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. Add the search 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.

Set up and manage indexing

Set up a search engine node and configure indexing.

Prerequisites

Ensure that:

- A vCenter datastore has been configured. Add a VMware vCenter Server on page 52 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: IPv4 only
 - 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 nodes in a multi-node search 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**—Note that only IPv4 addresses are supported.
 - **vCenter**—If you have added multiple vCenter Server instances, select the vCenter 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 backup data. 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.

- 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 search cluster.

NOTE: Ensure that the previous search engine node has successfully deployed before you add another node.

7. In the **Configure Search Engine** dialog box, enable or disable Search 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 supports the deletion of a Search Engine node from a multi-node Search cluster in the PowerProtect Data Manager UI.

You can delete an operational node from a Search cluster to decrease cluster capacity if the space is no longer required. You can also redeploy or delete nodes that could not be successfully added to the Search Engine.

When you delete an operational node, PowerProtect Data Manager moves the index data to the remaining nodes to avoid data loss.


When you delete a node, the operation is triggered and a new job is created, which you can view in the **Jobs > System Jobs** window to track its progress.

Delete operational nodes from a Search cluster

You can delete nodes that have been added to PowerProtect Data Manager and are in an operational state.

About this task

Before you can delete the primary node, you must delete all other nodes.

 **CAUTION:** Deleting the primary node deletes the index data and makes the Search cluster inactive. Add a node to make the Search cluster operational.

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 click **More Actions > Delete Node**.

In the **Delete Search Engine Node** window, choose one of the following options:

- Delete the node without data loss.

To delete the node and move the index data to the remaining nodes in the cluster, click **Delete Node**.

- Delete the node and its data.

 **CAUTION:** By selecting this option, the Search Engine deletes the node without redistributing the data to the remaining nodes in the cluster.

When you delete the node and the index data, the Search cluster becomes inactive.

To allow the Search Engine to delete the node along with the index data it holds, select the check box and click **Delete Node**.

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

Results

The node is deleted from the cluster.

Redeploy or delete failed nodes from a Search cluster

PowerProtect Data Manager enables you to redeploy or delete search nodes that could not be successfully deployed.

About this task

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

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Search Engine**.
2. Select the failed node that you want to either redeploy or delete from the Search cluster.
3. Do one of the following:

- To redeploy the failed node, click **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 for the node is correct.

- To delete the failed node, click **More Actions > Delete Node**.

Results

You can view the details for the operation in the **Jobs > System Jobs** window.

Next steps

Optionally, if you want to update the DNS and/or gateway during the search node redeployment, you can use one of the following commands:

- To update both the gateway and DNS, run `./infranodemgmt redeploy -node_id Search Node ID -updateDns DNS IPv4 address -updateGateway Gateway IPv4 address`
- To update the gateway only, run `./infranodemgmt redeploy -node_id Search Node ID -updateGateway Gateway IPv4 address`
- To update DNS only, run `./infranodemgmt redeploy -node_id Search Node ID -updateDns DNS IPv4 address`

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.

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.

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.

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**.

Perform a search

When the PowerProtect Search Engine is installed 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 virtual machine backup copies. When asset types are set up for index searching, the **File Search** button appears in the **Restore** menu for virtual machine assets.

Before performing a search, ensure that:

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

Virtual machine file level restore from a search

Within the **Restore** window of the PowerProtect Data Manager UI, **File Search** enables you to restore files from protected virtual machine backup copies to:

- The original virtual machine
- An alternate virtual machine.

NOTE: Only file level virtual machine restore is available from **File Search**.

File level restore to original virtual machine using File Search

Use **File Search** in the PowerProtect Data Manager UI to restore files from multiple copies across one or more virtual machines to the same location on the original vCenter Server. Only the Administrator and the Restore Administrator roles can restore data.

Prerequisites

- Review the section Supported platform versions for file-level restore for supported platform and operating system versions.
- Review the section File-level restore and SQL restore limitations on page 247.

NOTE: For file level restores to the original machine:

- The files must be restored from a Windows backup to a Windows machine, or from a Linux backup to a Linux machine.
- Restoring files from multiple copies with identical file names and paths from the same asset is not supported. In this case, only a file-level restore to the alternate virtual machine is available.

Steps

1. From the PowerProtect Data Manager UI, select **Restore > Assets**, and then select the **Virtual Machine** tab. The **Restore** window displays all the virtual machines available for restore.
2. Click **File Search**, and then perform the following:
 - a. Select a virtual machine from the **VM 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 files 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. Repeat steps b through d to select files from other virtual machines and copies. When finished with your selections, click **Restore**.

The **VM File Restore** wizard appears, displaying the **Location** page.

3. On the **Location** page:
 - a. Select **Restore to Original Location**.
 - b. Optionally, select **Overwrite existing files with the same name** to replace files in the original location with the files being restored if the files have the same name.
 - c. If you selected files from multiple virtual machines, and these virtual machines share the same credentials, move the **Use one set of credentials for all VMs** slider to the right to avoid retyping the credentials for each virtual machine.
 - d. For one or more virtual machines, type the virtual machine **User Name** and **Password**, and then click **Verify** to validate the credentials.
 - If there are administrator-level credentials that are associated with the virtual assets or protection policy being restored, specify end-user credentials.
 - If there are no administrator-level credentials that are associated with the virtual assets or protection policy being restored, specify administrator credentials. These credentials are handled as end-user credentials.

You are not required to wait for validation to complete before clicking **Verify** for another set of virtual machine credentials.

When validated, the **FLR Agent** is installed automatically on the restore destination, if it is not already installed. The **FLR Agent** facilitates the mounting and unmounting of disks and the browsing of files in the destination virtual machine and the backup copy. In order to complete the automatic **FLR Agent** installation, on Windows virtual machines the user must be an administrator account, and on Linux virtual machines the user must be the root user account, or a user in the operating system's local sudoers list. The section **FLR Agent for virtual machine file level restore** on page 248 provides more information.

- e. Optionally, leave **Keep FLR Agent Installed** selected if you do not want to remove the **FLR Agent** on the destination virtual machines after the restore completes.
 - f. Click **Next**.
- The **Summary** page appears.
4. On the **Summary** page:
 - a. Review the information to ensure that the restore details are correct. You can click **Edit** next to certain rows to change the information.
 - b. Click **Restore** or **Finish**.

5. Go to the **Jobs** window to monitor the restore. A batch file level restore job with multiple files 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 virtual machine using File Search

Use **File Search** in the PowerProtect Data Manager UI to restore files from multiple copies across one or more virtual machines to a new location on a new virtual machine. This restore can be performed to the primary vCenter (the location of the original virtual machine), or a secondary vCenter Server. Only the Administrator and the Restore Administrator roles can restore data.

Prerequisites

- Review the section **Supported platform versions for file-level restore** for supported platform and operating system versions.
- Review the section **File-level restore and SQL restore limitations** on page 247.

NOTE: For file-level restores to an alternate virtual machine:

- You can only restore files from a Windows backup to a Windows machine, or from a Linux backup to a Linux machine.
- Restore of multiple files from different operating systems to the same target virtual machine is not supported. In this case, only a file level restore to the original virtual machine is available.

Steps

1. From the PowerProtect Data Manager UI, select **Restore > Assets**, and then select the **Virtual Machine** tab.
The **Restore** window displays all the virtual machines available for restore.
2. Click **File Search**, and then perform the following:
 - a. Select a vCenter from the **vCenter Name** list.
 - b. Select a virtual machine from the **VM Name** list.
 - c. 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.
 - d. In the **Results** pane, select the files 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.
 - e. Repeat steps b through d to select files from other virtual machines and copies. When finished with your selections, click **Restore**.
The **VM File Restore** wizard appears, displaying the **Location** page.
3. On the **Location** page:
 - a. Select **Restore to Alternate Location**.
The table on the page updates to display the available destination virtual machines within the vCenter and, when a specific virtual machine is selected, its location.
 - b. Expand the vCenter to locate the virtual machine that you want to restore to, and then select the virtual machine.
A prompt appears, requesting the credentials of this virtual machine.
 - c. Type the virtual machine **User Name** and **Password**, and then click **Verify** to validate the credentials.
When validated, the **FLR Agent** is installed automatically on the restore destination, if it is not already installed. The **FLR Agent** facilitates the mounting and unmounting of disks and the browsing of files in the destination virtual machine and the backup copy. In order to complete the automatic **FLR Agent** installation, on Windows virtual machines the user must be an administrator account, and on Linux virtual machines the user must be the root user account, or a user in the operating system's local sudoers list. The section **FLR Agent for virtual machine file level restore** on page 248 provides more information.
 - d. Optionally, leave **Keep FLR Agent installed** selected if you do not want to remove the **FLR Agent** on the destination virtual machines after the restore completes.
 - e. When validation completes, click **Close** to return to the **Location** page.
The **Location** page updates with the available destination folders on the selected virtual machine.
 - f. Browse to the destination folder, or select a location and click **Add Folder** to create a destination within this folder.
 - g. Optionally, select **Overwrite existing files with the same name** to replace files in the destination folder with the files being restored if the files have the same name.
 - h. Click **Next**.
The **Summary** page appears.
4. On the **Summary** page:
 - a. Review the information to ensure that the restore details are correct. You can click **Edit** next to certain rows to change the information. If you are not restoring to the original virtual machine, an additional field appears for the **Target VM**.
 - b. Click **Restore** or **Finish**.
5. Go to the **Jobs** window to monitor the restore.
A batch file level restore job with multiple files 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.

Troubleshooting Search Engine issues

This section lists troubleshooting and Search Engine issues.

Error displays during node failure

The following error might display during a search when a 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 node, and then retry the operation. If you choose to edit, delete the node and the new mode modal appears with your previous input. 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 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-d5ec56b9-69ec-4183-9c94-7c0230408765"
```

- Examine the rest-engine logs in the search 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 in the node(s) to investigate further, if necessary, regenerate the certificate files.

Access the Search Node to discover passwords

Use the following steps to discover the admin and root passwords for all deployed search nodes:

1. Connect to the PowerProtect Data Manager console and change to the root user.
2. Change directory to /opt/emc/vmdirect.
3. Source unit/vmdirect.env.
4. Run bin/infranodemgmt get -secret.

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 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 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-a9a7-6080649d0d24

error 18 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 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 node FQDN.
3. Run `/usr/local/brs/puppet/scripts/generate_certificates.sh -n -c -b <search node FQDN>`
A properties file is created in the `/root` directory called `<search node FQDN>.properties`.
4. Open this file to determine the location of the generated certificates. They should be located in `/etc/ssl/certificates/<search node FQDN>`.
5. From a separate terminal, SSH into the search node using the password that was revealed with the `infranodemgmt Get` call in step 2.
6. Change directory to `/var/lib/dellemc/vmboot/trust` and move the `key`, `cert`, and `thumbprint` files over.
7. Copy the certificate files that were generated in PowerProtect Data Manager as follows:
 - `otca.pem` to `thumbprint`
 - `<search node FQDN>key.pem` to `key`
 - `<search node FQDN>.pem` to `cert`
8. Paste the files to `/var/lib/dellemc/vmboot/trust`.
9. Set the permissions for the `key`, `cert`, and `thumbprint` files to `0644`, and then set the ownership of these files to `root:app`
10. Restart the rest-engine daemon or the vproxyd daemon) to pick up the new certificates: `systemctl restart search-rest-engine`.
11. 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 search service is functional.

Search cluster is full

If the search cluster is full, you can deploy additional nodes by following the steps in [Set up and manage indexing](#) on page 46.

If the search cluster 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 Cluster.

To remove indexes manually, complete the following steps:

1. Use SSH to log in to the Search virtual machine.
2. Create a snapshot of the Search 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, ./searchmgmt -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 nodes on the PowerProtect Data Manager Search cluster are configured with IP addresses that can be accessed externally. These nodes are configured with admin or root user accounts, which are only used to log in to the Search nodes for troubleshooting software issues. The password management policies for these accounts are set to lock the admin user account if there are three wrong password attempts within a 5 minute time period. If you try to access the node while the admin user account is locked, the amount of time that the account remains locked increases.

There is no public interface available that enables you to access the search node by using admin credentials. All required information about the Search Engine nodes is obtained through the PowerProtect Data Manager UI.

A Search node might become locked for the following reasons:

- A user or program tries to SSH into the search node and makes three wrong attempts at entering the password.
- Running monitoring software that tries to log in to the Search node with the wrong admin credentials and locks the system.
- Running Penetration Testing (PEN) on the VMs in the vCenter.

The Search admin user account enables the PowerProtect Data Manager system to perform different operations on the Search 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 Search cluster is in a failed state, the entire Search cluster becomes unavailable. As a result, the Search cluster is unable to perform any indexing or search operations.

Workaround

To work around this issue, first access the Search node to discover the admin and root credentials for the node. After you discover the node credentials, log in to the node through the vCenter console to reset the admin credentials.

Use the following steps to discover the admin and root passwords for all deployed Search nodes:

1. Connect to the PowerProtect Data Manager console and change to the root user.
2. Change directory to `/opt/emc/vmdirect`.
3. Run `unit/vmdirect.env`
4. Run `bin/infranodemgmt get -secret`

Before you access the Search node through the vCenter console, determine why the admin account is locked.

Use the following steps to unlock the admin user account:

1. Log in to the vCenter where the Search node is present.
2. Select the Search node from the **VMs and Templates** view in the left pane of the vSphere Client home page.
3. Launch a PowerProtect Data Manager VM vCenter console.
4. Log in to the vCenter console with the root username and password.
5. Run the following command to reset the admin user account credentials:

```
/sbin/pam_tally2 --user admin --reset
```

Managing Assets

Topics:

- About asset sources, assets, and storage
- Prerequisites for discovering asset sources
- Enable an asset source
- Delete an asset source
- Adding a vCenter Server asset source
- VM Direct protection engine overview
- 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 Dell EMC PowerProtect DD Management Center (DDMC) as the storage and programmatic interface for controlling protection storage systems.

Asset sources can be a vCenter Server, Kubernetes cluster, application host, SMIS server, or Cloud Snapshot Manager tenant. Assets can be virtual machines, Exchange databases, SQL databases, Oracle databases, SAP HANA databases, file systems, Kubernetes namespaces, or storage groups.

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

About vCenter Server asset sources and virtual assets

After you add a vCenter Server as an asset source in PowerProtect Data Manager, an automatic discovery of VMware entity information from the vCenter Server is initiated.

The virtual assets for the vCenter Server appear in the **Assets** window of the PowerProtect Data Manager UI under the **Virtual Machine** tab.

The initial vCenter Server discovery identifies all ESXi clusters, hosts, and virtual machines within the vCenter Server. Subsequent discoveries can be performed to identify any additional or changed VMware entities since the last discovery operation. You can also manually initiate a discovery of VMware entities at any time from the **vCenter** tab of the **Asset Sources** window by selecting a vCenter Server and clicking **Discover**.

Upon vCenter Server and virtual asset discovery, the PowerProtect Data Manager VM Direct protection engine facilitates the management of virtual assets as PowerProtect Data Manager resources for the purposes of backup and recovery. Dell EMC recommends that you also add an external VM Direct Engine in the **Protection Engines** window. You can protect virtual machine assets by manually adding the assets to a virtual machine protection policy, or by creating and applying protection rules to determine which assets are included in a protection policy based on rule definitions.

About other asset sources

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

NOTE: The *PowerProtect Data Manager Administration and User 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.

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 Exchange agent

After the Microsoft Exchange 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 Exchange application data on the application host, and to check and monitor backup compliance against protection policies.

Microsoft SQL agent

After the Microsoft SQL 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 SQL 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.

Prerequisites for discovering asset sources

Perform these tasks before you discover the asset sources.

- 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 on page 60 provides more information.
- Configure all asset sources with an NTP server.
- Before you register an SQL application, ensure that the DD system has been discovered successfully.
- For discovery of Application Agent and File System asset sources:
 - Ensure that all clocks on both the App/File System host and PowerProtect Data Manager are time-synced to the local NTP server to ensure discovery of the backups.
 - Ensure that the App/File System host and the PowerProtect Data Manager network can see/resolve each other.
 - Ensure that port 7000 is open on the App/File System host.
- Discovery of a vCenter Server asset source will exclude the following:
 - Virtual machines with a status of **Inaccessible**, **Invalid**, or **Orphaned**.
 - The virtual machine template
 - The shadow (or standby) virtual machine created by Dell EMC RecoverPoint for Virtual Machines, also referred to as the vRPA copy.

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

Enable an asset source

An asset source, such as a vCenter Server, 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.

There are some 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 UI, 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 UI. On a new installation, however, no asset sources are enabled by default.

Steps

1. From the PowerProtect Data Manager UI, 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 agent, 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 virtual machines, then you cannot disable the vCenter asset source. But if you register a vCenter Server and then delete the vCenter 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 can be deleted automatically or manually. The *PowerProtect Data Manager Administration and User 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 vCenter Server asset source

After you register a vCenter Server with PowerProtect Data Manager, you can use the **Asset Sources** window in the PowerProtect Data Manager UI to add a vCenter Server asset source to the PowerProtect Data Manager environment.

Adding a vCenter Server asset source is required if you want to schedule a backup through PowerProtect Data Manager.

Add a VMware vCenter Server

Perform the following steps to add a vCenter Server as an asset source in the PowerProtect Data Manager UI for virtual machine protection and Tanzu Kubernetes guest cluster protection.

Prerequisites

- Ensure that the asset source is enabled. Enable an asset source on page 60 provides instructions.
- Log in as a user with the Administrator role. Only the Administrator role can manage asset sources.
- By default, PowerProtect Data Manager enforces SSL certificates during communication with vCenter Server. If a certificate appears and you trust the certificate, click **Verify**.

Note, however, that SSL certificate enforcement requires that the common name (cn) of the x509 certificate on the vCenter Server matches the hostname of the vCenter URL. The common name of the x509 certificate is typically the vCenter server fully qualified domain name (FQDN), but it could be the vCenter server IP address. You can inspect the vCenter server SSL certificate to determine whether the x509 common name is an FQDN or IP. When creating an asset source resource, in order to pass SSL certificate enforcement, the asset source resource hostname must match the common name of the x509 certificate on the vCenter server.

NOTE: It is highly recommended that you do not disable certificate enforcement. If disabling the certificate is required, carefully review the instructions in the section [Disable vCenter SSL certificate validation](#) on page 246.

Steps

1. From the left navigation pane, select **Infrastructure > Asset Sources**.
The **Asset Sources** window appears.
2. Select the **vCenter** tab.
3. Click **Add**.
The **Add vCenter** dialog displays.
4. Specify the source attributes:
 - a. In the **Name** field, specify the vCenter Server name.
 - b. In the **Address** field, specify the fully qualified domain name (FQDN) or the IP address.
NOTE: For a vCenter Server, it is recommended that you use the FQDN instead of the IP address.
 - c. In the **Port** field, specify the port for communication if you are not using the default port, 443.
5. Under **Host Credentials**, choose an existing entry from the list to use for the vCenter user credentials. Alternatively, you can click **Add** from this list to add new credentials, and then click **Save**.
NOTE: Ensure that you specify the credentials for a user whose role is defined at the vCenter level, as opposed to being restricted to a lower-level container object in the vSphere object hierarchy.
6. If you want to make a subset of the PowerProtect Data Manager UI functionality available within the **vSphere Client**, move the **vSphere Plugin** slider to the right.
Available functionality includes:
 - The monitoring of active virtual machine/VMDK protection policies, and
 - Restore options such as **Restore to Original**, **Restore to New**, and **Instant Access**.**NOTE:** You can unregister the vSphere plug-in at any time by moving the slider to the left.
7. By default, the vCenter discovery occurs automatically after adding the vCenter, and subsequent discoveries are incremental. If you want to schedule a full discovery at a certain time every day, move the **Schedule Discovery** slider to the right, and then specify a time.
8. If there is no hosting vCenter and you want to make this the vCenter Server that hosts PowerProtect Data Manager, select **Add as hosting vCenter**. If a vCenter Server has already been added as the hosting vCenter, this option will be greyed out. Specify a vCenter Server as the PowerProtect Data Manager host on page 157 provides more information about adding a host vCenter.
9. If the vCenter server SSL certificate cannot be trusted automatically, a dialog box appears requesting certificate approval. Review the certificate, and then click **Verify**.
10. Click **Save**.

The vCenter Server information that you entered now appears as an entry in a table on the **Asset Sources** window. You can click the magnifying glass icon next to the entry to view more details, such as the next scheduled discovery, the number of assets within the vCenter, and whether the **vSphere Plugin** is enabled.

NOTE: Although PowerProtect Data Manager automatically synchronizes with the vCenter server under most circumstances, certain conditions might require you to initiate a manual discovery.

After discovery, PowerProtect Data Manager starts an incremental discovery in the background periodically to keep updating PowerProtect Data Manager with vCenter changes. You can always do an on-demand discovery.

11. Optionally, you can set warning and failure thresholds for the available space on the datastore. Setting these thresholds enables you to check if enough storage space is available in the datastore to save the snapshot of the virtual machine during the backup process. The backup completes with a warning in the logs if the available free space in the datastore is less than or equal to the percentage indicated in the **Datastore Free Space Warning Threshold**. The backup fails if the available free space in the datastore is less than or equal to the percentage indicated in the **Datastore Free Space Failure Threshold**. To add Datastore Free Space Warning and Failure Thresholds:

- a. Click the gear icon to open the **vCenter Settings** dialog.
- b. Type a percentage value to indicate when a warning message should display due to low datastore free space.
- c. Type a percentage value to indicate when a virtual machine backup failure should occur due to low datastore free space.
- d. Click **Save**.

NOTE: Datastore free space thresholds are disabled by default.

12. Select **Infrastructure > Assets**.

The **Assets** window appears.

13. If not already selected, click the **Virtual Machine** tab.

Results

Upon a successful discovery of the vCenter server asset source, the virtual machine assets in the vCenter display in the **Infrastructure > Assets** window.

You can modify the details for the vCenter asset source by selecting the vCenter in the **Infrastructure > Asset Sources** window and clicking **Edit**. You cannot, however, clear the **Add as hosting vCenter** checkbox when editing an asset source if this vCenter Server has already been added as the hosting vCenter. For this operation, use the **Hosting vCenter** window, as described in the section *Specify a vCenter Server as the PowerProtect Data Manager host* on page 157.

NOTE: Discovery time is based on networking bandwidth. The resources that are discovered and the resources that are performing the discovery impact performance each time that you initiate a discovery process. It might appear that PowerProtect Data Manager is not updating the Asset Sources data while the discovery is in progress.

Next steps

Add a VM Direct appliance to facilitate data movement, and then create virtual machine protection policies to back up these assets. The PowerProtect Data Manager software comes bundled with an embedded VM Direct engine, which is automatically used as a fallback proxy for performing backups and restores when the added external proxies fail or are disabled. It is recommended that external proxies be deployed since the embedded VM Direct engine has limited capacity for performing backup streams. To add a VM Direct Engine, select **Infrastructure > Protection Engines**.

Creating a dedicated vCenter user account

Dell EMC strongly recommends that you set up a separate vCenter user account at the root level of the vCenter that is strictly dedicated for use with PowerProtect Data Manager and the VM Direct protection engine.

Use of a generic user account such as "Administrator" could make future troubleshooting efforts difficult as it might not be clear which "Administrator" actions are actually interfacing or communicating with PowerProtect Data Manager. Using a separate vCenter user account ensures maximum clarity if it becomes necessary to examine vCenter logs.

You can specify the credentials for a vCenter user account when you add the vCenter as an asset source in the UI. When you add the vCenter, ensure that you specify a user whose role is defined at the vCenter level and not restricted to a lower level container object in the vSphere object hierarchy.

Specify the required privileges for a dedicated vCenter user account

You can use the **vSphere Client** to specify the required privileges for the dedicated vCenter user account, or you can use the **PowerCLI**, which is an interface for managing vSphere.

The following table includes the privileges required for this user:

NOTE: For the privileges required when administering PowerProtect Data Manager in a cloud environment, see Specify the required privileges for a dedicated cloud-based vCenter user account on page 193. For the additional privileges required when using the Transparent Snapshot Data Mover (TSDM) protection mechanism for virtual machine crash-consistent data protection, see Additional privileges required for a dedicated vCenter user account to use Transparent Snapshot Data Mover on page 178.

Table 25. Minimum required vCenter user account privileges

Setting	vCenter 6.5 and later required privileges	PowerCLI equivalent required privileges
Alarms	<ul style="list-style-type: none"> Create alarm Modify alarm 	<pre>\$privileges = @('System.Anonymous', 'System.View', 'System.Read', 'Alarm.Create', 'Alarm.Edit', 'Cryptographer.AddDisk', 'Cryptographer.Access', 'Cryptographer.RegisterVM', 'Datastore.Rename', 'Datastore.Move', 'Datastore.Delete', 'Datastore.Browse', 'Datastore.DeleteFile', 'Datastore.FileManagement', 'Datastore.AllocateSpace', 'Datastore.Config', 'Extension.Register', 'Extension.Unregister', 'Extension.Update', 'Folder.Create', 'Global.ManageCustomFields', 'Global.SetCustomField', 'Global.LogEvent', 'Global.CancelTask', 'Global.Licenses', 'Global.Settings', 'Global.DisableMethods', 'Global.EnableMethods', 'Host.Config.Storage', 'InventoryService.Tagging.AttachTag', 'InventoryService.Tagging.ObjectAttachable', 'InventoryService.Tagging.CreateTag', 'InventoryService.Tagging.CreateCategory', 'Network.Config', 'Network.Assign', 'Resource.AssignVMtoPool', 'Resource.HotMigrate', 'Resource.ColdMigrate', 'Sessions.ValidateSession', 'StorageProfile.Update', 'StorageProfile.View', 'Task.Create', 'Task.Update', 'VApp.ApplicationConfig', 'VApp.Export', 'VApp.Import', 'VirtualMachine.Config.Rename', 'VirtualMachine.Config.Annotation', 'VirtualMachine.Config.AddExistingDisk'</pre>
Cryptographic operations	<ul style="list-style-type: none"> Add disk Direct Access Register VM 	
Datastore	<ul style="list-style-type: none"> Allocate space Browse datastore Configure datastore Low level file operations Move datastore Remove datastore Remove file Rename datastore 	
Extension	<ul style="list-style-type: none"> Register extension Unregister extension Update extension 	
Folder	<ul style="list-style-type: none"> Create folder 	
Global	<ul style="list-style-type: none"> Cancel task Disable methods Enable methods Licenses Log event Manage custom attributes Set custom attribute Settings 	
Host	<ul style="list-style-type: none"> Configuration > Storage partition configuration 	
vSphere Tagging	<ul style="list-style-type: none"> Assign or Unassign vSphere Tag Assign or Unassign vSphere Tag on Object <p>NOTE: This only applies to vCenter 7.0 and later.</p> <ul style="list-style-type: none"> Create vSphere Tag Create vSphere Tag Category 	
Network	<ul style="list-style-type: none"> Assign network Configure 	
Resource	<ul style="list-style-type: none"> Assign virtual machine to resource pool Migrate powered off virtual machine Migrate powered on virtual machine 	

Table 25. Minimum required vCenter user account privileges (continued)

Setting	vCenter 6.5 and later required privileges	PowerCLI equivalent required privileges
Sessions	<ul style="list-style-type: none"> Validate session 	
SPBM policy restore	<ul style="list-style-type: none"> Profile-driven storage Profile-driven storage update Profile-driven storage view 	'VirtualMachine.Config.AddNewDisk', 'VirtualMachine.Config.RemoveDisk', 'VirtualMachine.Config.RawDevice', 'VirtualMachine.Config.HostUSBDevice', 'VirtualMachine.Config.CPUCount', 'VirtualMachine.Config.Memory', 'VirtualMachine.Config.AddRemoveDevice'
Tasks	<ul style="list-style-type: none"> Create task Update task 	
vApp	<ul style="list-style-type: none"> Export Import vApp application configuration 	'VirtualMachine.Config.EditDevice', 'VirtualMachine.Config.Settings', 'VirtualMachine.Config.Resource', 'VirtualMachine.Config.UpgradeVirtualHardware', 'VirtualMachine.Config.ResetGuestInfo', 'VirtualMachine.Config.AdvancedConfig', 'VirtualMachine.Config.DiskLease', 'VirtualMachine.Config.SwapPlacement', 'VirtualMachine.Config.DiskExtend', 'VirtualMachine.Config.ChangeTracking', 'VirtualMachine.Config.ReloadFromPath', 'VirtualMachine.Config.ManagedBy', 'VirtualMachine.GuestOperations.Query', 'VirtualMachine.GuestOperations.Modify'
Virtual Machine		
Change Configuration	<ul style="list-style-type: none"> Acquire disk lease Add existing disk Add new disk Add or remove device Advanced configuration Change CPU count Change Memory Change Settings Change Swapfile placement Change resource Configure Host USB device Configure Raw device Configure managedby Extend virtual disk Modify device settings Reload from path Remove disk Rename Reset guest information Set annotation Toggle disk change tracking Upgrade virtual machine compatibility 	'VirtualMachine.GuestOperations.Execute', 'VirtualMachine.Interact.PowerOn', 'VirtualMachine.Interact.PowerOff', 'VirtualMachine.Interact.Reset', 'VirtualMachine.Interact.ConsoleInteraction', 'VirtualMachine.Interact.DeviceConnection', 'VirtualMachine.Interact.SetCDMedia', 'VirtualMachine.Interact.ToolsInstall', 'VirtualMachine.Interact.GuestControl', 'VirtualMachine.Inventory.Create', 'VirtualMachine.Inventory.Register', 'VirtualMachine.Inventory.Delete', 'VirtualMachine.Inventory.Unregister', 'VirtualMachine.Provisioning.DiskRandomAccess', 'VirtualMachine.Provisioning.DiskRandomRead', 'VirtualMachine.Provisioning.GetVmFiles', 'VirtualMachine.Provisioning.MarkAsTemplate', 'VirtualMachine.State.CreateSnapshot', 'VirtualMachine.State.RevertToSnapshot'
Edit Inventory	<ul style="list-style-type: none"> Create new Register Remove Unregister 	'VirtualMachine.State.RemoveSnapshot', 'VirtualMachine.State.RemoveSnapshot', }
Guest operations	<ul style="list-style-type: none"> Guest operation modifications Guest operation program execution Guest operation queries 	
Interaction	<ul style="list-style-type: none"> Configure CD media Connect devices Console interaction Guest operating system management by VIX API Install VMware Tools Power off Power on Reset 	New-VIRole -Name 'PowerProtect' -Privilege (Get-VIPrivilege -Id \$privileges)
Provisioning	<ul style="list-style-type: none"> Allow disk access 	

Table 25. Minimum required vCenter user account privileges (continued)

Setting	vCenter 6.5 and later required privileges	PowerCLI equivalent required privileges
	<ul style="list-style-type: none"> Allow read-only disk access Allow virtual machine download Mark as template 	
Snapshot Management	<ul style="list-style-type: none"> Create snapshot Remove snapshot Revert to snapshot 	

VM Direct protection engine overview

The VM Direct protection engine provides two functions within PowerProtect Data Manager:

- A virtual machine data protection solution—Deploy a VM Direct Engine in the vSphere environment to perform virtual machine snapshot backups, which improves performance and reduces network bandwidth utilization by using the protection storage source-side deduplication.
- A Tanzu Kubernetes guest cluster data protection solution—Deploy a VM Direct Engine in the vSphere environment for protection of vSphere CSI-based persistent volumes, for which it is required to use a VM Proxy instead of the cProxy, for the management and transfer of backup data.

The VM Direct protection engine is enabled after you add a vCenter Server in the **Asset Sources** window, and allows you to collect VMware entity information from the vCenter Server and save VMware virtual machines and Tanzu Kubernetes guest cluster namespaces and PVCs as PowerProtect Data Manager resources for the purposes of backup and recovery.

To view statistics for the VM Direct engine, manage and monitor VM Direct appliances, and add an external VM Direct appliance to facilitate data movement, select **Infrastructure > Protection Engines**. Add a VM Direct Engine on page 66 provides more information.

NOTE: In the **VM Direct Engines** pane, **VMs Protected** refers to the number of assets protected by PowerProtect Data Manager. This count does not indicate that all of the virtual machines have been protected successfully. To determine the success or failure of asset protection, use the **Jobs** window.

When you add an external VM Direct appliance, the **VM Direct Engines** pane provides the following information:

- The VM Direct appliance IP address, name, gateway, DNS, network, and build version. This information is useful for troubleshooting network issues.
- The vCenter and ESXI hostname.
- The VM Direct appliance status (green check mark if the VM Direct appliance is ready, red x if the appliance is not fully operational). The status includes a short explanation to help you troubleshoot the VM Direct Engine if the VM Direct appliance is not in a fully operational state.
- The transport mode that you selected when adding the VM Direct appliance (Hot Add, Network Block Device, or the default setting Hot Add, Fallback to Network Block Device).

Requirements for an external VM Direct Engine

When adding an external VM Direct Engine, note the following system requirements:

- CPU: 4 * 2 GHz (4 virtual sockets, 1 core for each socket)
- Memory: 8 GB RAM
- Disks: 2 disks (59 GB and 98 GB)
- Internet Protocol: IPv4 only
- SCSI controller: maximum of 4
- NIC: One vmxnet3 NIC with one port

Add a VM Direct Engine

Perform the following steps in the **Protection Engines** window of the PowerProtect Data Manager UI to deploy an external VM Direct Engine, also referred to as a VM proxy. The VM Direct engine facilitates data movement for virtual machine.

protection policies, Kubernetes cluster protection policies that require a VM proxy instead of the cProxy, and network attached storage (NAS) protection policies.

Prerequisites

Review the sections [Requirements for an external VM Direct Engine on page 66](#) and [Transport mode considerations on page 252](#).

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

About this task

The PowerProtect Data Manager software comes bundled with an embedded VM Direct engine, which is automatically used as a fallback proxy for performing backups and restores when the added external proxies fail or are disabled. Dell Technologies recommends that you deploy external proxies by adding a VM Direct Engine for the following reasons:

- An external VM Direct Engine for VM proxy backup and recovery can provide improved performance and reduce network bandwidth utilization by using source-side deduplication.
- The embedded VM Direct engine has limited capacity for backup streams.
- The embedded VM Direct engine is not supported for VMware Cloud on AWS operations.

An external VM Direct engine is not required for virtual machine protection policies that use the Transparent Snapshot Data Mover (TSDM) protection mechanism. For these policies, the embedded VM Direct engine is sufficient.

NOTE: Cloud-based OVA deployments of PowerProtect Data Manager do not support the configuration of data-traffic routing or VLANs. Those deployments skip the **Networks Configuration** page.

Steps

1. From the left navigation pane, select **Infrastructure > Protection Engines**.
The **Protection Engines** window appears.
2. In the **VM Direct Engines** pane of the **Protection Engines** window, click **Add**.
The **Add Protection Engine** wizard displays.
3. On the **Protection Engine Configuration** page, complete the required fields, which are marked with an asterisk.
 - **Hostname, Gateway, IP Address, Netmask, and Primary DNS**—Note that only IPv4 addresses are supported.
 - **vCenter to Deploy**—If you have added multiple vCenter Server instances, select the vCenter on which to deploy the protection engine.
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 protection engine.
 - **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 based on ranking (whether the datastores are shared or local), and available capacity (the datastore with the most capacity appearing at the top of the list).

You can choose the specific datastore on which the protection engine resides, or leave the default selection of **automatic** to allow PowerProtect Data Manager to determine the best location to host the protection engine.
 - **Transport Mode**—Select **Hot Add**.
 - **Supported Protection Type**—Select whether this protection engine is intended for **Virtual Machine, Kubernetes Tanzu guest cluster, or NAS asset protection**.
4. Click **Next**.
5. On the **Networks Configuration** page:
If this is a cloud-based OVA deployment of PowerProtect Data Manager, click **Next** and proceed to step 7.

The **Networks Configuration** page configures the virtual network (VLAN) to use for backup data. 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 VST (Virtual Switch Tagging) or VGT (Virtual Guest Tagging) network.
If you select a VGT portgroup, the list displays all virtual networks within the trunk range. If you select a VST portgroup, the list displays only the virtual network for the current VLAN ID.
 - b. Select one or more virtual networks from the list.

A protection engine 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.

- c. 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 selected network. Verify the proposed IP addresses.

- d. Click **Next**.

6. When adding a VM Direct engine for Kubernetes guest cluster protection, add a second network interface card (NIC) if the PowerProtect controller pod running in the guest cluster cannot reach the vProxy on the primary network. Provide information for the second NIC, and then click **Next**.

7. On the **Summary** page, review the information and then click **Finish**.

The protection engine is added to the **VM Direct Engines** pane. An additional column indicates the engine purpose. Note that it can take several minutes to register the new protection engine in PowerProtect Data Manager. The protection engine also appears in the **vSphere Client**.

Results

When an external VM Direct Engine is deployed and registered, PowerProtect Data Manager uses this engine instead of the embedded VM Direct engine for any data protection operations that involve virtual machine protection policies. If all external VM Direct Engines are unavailable, PowerProtect Data Manager uses the embedded VM Direct engine as a fallback to perform limited scale backups and restores. If you do not want to use the external VM Direct Engine, you can disable this engine. Additional VM Direct actions on page 68 provides more information.

- NOTE:** The external VM Direct Engine is always required for VMware Cloud on AWS operations, Kubernetes cluster protection policies that require a VM Proxy instead of the cProxy, and NAS protection policies. If no external VM Direct Engine is available for these solutions, data protection operations fail.

Next steps

If the protection engine deployment fails, review the network configuration of PowerProtect Data Manager in the **System Settings** window to correct any inconsistencies in network properties. After successfully completing the network reconfiguration, delete the failed protection engine and then add the protection engine in the **Protection Engines** window.

When configuring the VM Direct Engine in a VMware Cloud on AWS environment, if you deploy the VM Direct Engine to the root of the cluster instead of inside the Compute-ResourcePool, you must move the VM Direct Engine inside the Compute-ResourcePool.

Additional VM Direct actions

For additional VM Direct actions, such as enabling, disabling, redeploying, or deleting the VM Direct Engine, or changing the network configuration, use the **Protection Engines** window in the PowerProtect Data Manager UI. To throttle the capacity of VM Direct engines, use a command-line tool on PowerProtect Data Manager.

To get external VM Direct Engine credentials, see the procedure in the *PowerProtect Data Manager Security Configuration Guide*.

Disable a VM Direct Engine

You can disable an added VM Direct Engine that you do not currently require for virtual machine backup and recovery. To disable a VM Direct Engine:

1. On the **Protection Engines** window, select the VM Direct Engine that you want to disable from the table in the **VM Direct Engines** pane.
2. In the far right of the **VM Direct Engines** pane, click the three vertical dots.
3. From the menu, select **Disable**.

NOTE: A disabled VM Direct Engine is not used for any new protection activities, and is not automatically updated during a PowerProtect Data Manager update.

Delete a VM Direct Engine

When you disable a VM Direct Engine, the **Delete** button is enabled. If you no longer require the VM Direct Engine, perform the following steps to delete the engine:

1. On the **Protection Engines** window, select the VM Direct Engine that you want to remove from the table in the **VM Direct Engines** pane.
2. In the far right of the **VM Direct Engines** pane, click the three vertical dots.
3. From the menu, select **Disable**.
4. Click **Delete**.

Enable a disabled VM Direct Engine

When you want to make a disabled VM Direct Engine available again for running new protection activities, perform the following steps to re-enable the VM Direct Engine.

1. On the **Protection Engines** window, select the VM Direct Engine that you want to re-enable from the table in the **VM Direct Engines** pane.
2. In the far right of the **VM Direct Engines** pane, click the three vertical dots.
3. From the menu, select **Enable**.

NOTE: If a PowerProtect Data Manager version update occurred while the VM Direct Engine was disabled, a manual redeployment of the VM Direct Engine is also required.

Redeploy a VM Direct Engine

If a PowerProtect Data Manager software update occurred while a VM Direct Engine was disabled, or an automatic update of the VM Direct Engine did not occur due to network inaccessibility or an environment error, the **Redeploy** option enables you to manually update the VM Direct Engine to the version currently in use with the PowerProtect Data Manager software. Perform the following steps to manually redeploy the VM Direct Engine.

1. On the **Protection Engines** window, select the VM Direct Engine that you want to redeploy from the table in the **VM Direct Engines** pane.
2. In the far right of the **VM Direct Engines** pane, click the three vertical dots.
3. If the VM Direct Engine is not yet enabled, select **Enable** from the menu.
4. When the VM Direct Engine is enabled, select **Redeploy** from the menu.

The VM Direct Engine is redeployed with its previous configuration details.

Update the DNS or gateway during redeployment

Optionally, if you want to update the vProxy DNS and/or gateway during the VM Direct Engine redeployment, you can use one of the following commands:

- To update both the gateway and DNS, run `./vproxymgmt redeploy -vproxy_id VM Direct Engine ID -updateDns DNS IPv4 address -updateGateway Gateway IPv4 address`
- To update the gateway only, run `./vproxymgmt redeploy -vproxy_id VM Direct Engine ID -updateGateway Gateway IPv4 address`
- To update DNS only, run `./vproxymgmt redeploy -vproxy_id VM Direct Engine ID -updateDns DNS IPv4 address`

Edit the network configuration for a VM Direct Engine

If VM Direct Engine deployment failed because of a virtual network configuration problem, you can update the configuration to add additional IP addresses to the static IP pool. You can also add the VM Direct Engine to a virtual network in the same VGT port group.

Perform the following steps to change the network configuration:

1. On the **Protection Engines** window, select the VM Direct Engine from the table in the **VM Direct Engines** pane.
2. Click **Edit**.
3. Select the row that corresponds to the virtual network with the configuration error, or the virtual network to which you want to add the VM Direct Engine.
4. Type an available static IP address or IP address range in the **Additional IP Addresses** column.
5. Click **Next**.
6. On the **Summary** page, verify the network settings, and then click **Next**.

To change other network configuration settings, delete the VM Direct Engine and then deploy a new VM Direct Engine.

Throttle capacity of a VM Direct Engine

In performance-limited environments, you can use a command-line tool on PowerProtect Data Manager to reduce the maximum capacity of VM Direct engines.

- The default value for *VM Configured Capacity Units* of an external VM Direct engine is 100. The minimum value is 4.
- A VM Direct engine can backup one disk with 4 units of capacity at a time.

Perform these steps to throttle the capacity of one or more VM Direct engines:

1. Connect to the PowerProtect Data Manager console and change to the root user.
2. Type: `source /opt/emc/vmdirect/unit/vmdirect.env`
3. To view the list of VM Direct engines and their IDs, type: `/opt/emc/vmdirect/bin/vproxymgmt get -list`
4. To change the capacity of a VM Direct engine, type (once per engine): `/opt/emc/vmdirect/bin/vproxymgmt modify -vproxy_id [VProxy ID] -capacity [percentage]`
5. To verify the change in *VM Configured Capacity Units*, type: `/opt/emc/vmdirect/bin/vproxymgmt get -list`

Transparent Snapshot Data Mover protection mechanism

The **Protection Engines** window in the PowerProtect Data Manager UI includes a pane for **Transparent Snapshot Data Movers**. Transparent Snapshot Data Mover (TSDM) is a protection mechanism that is introduced in PowerProtect Data Manager 19.9 for data movement during virtual machine protection operations. Previously, the only protection mechanism available in PowerProtect Data Manager for virtual machine protection was the VMware vStorage API for Data Protection (VADP).

A vSphere Installation Bundle (VIB) is included with the software installation and update packages for PowerProtect Data Manager 19.9 to facilitate the use of TSDM, and is enabled at the vCenter level upon the PowerProtect Data Manager 19.9 installation or update. The VIB installation occurs automatically at the cluster level when a virtual machine protection policy is created, with no requirement to restart the ESXi hosts or put the hosts into maintenance mode. Any new virtual machine protection policies use TSDM as the default protection mechanism instead of VADP, provided that the vCenter/ESXi Server that hosts the virtual machines is a minimum version of 7.0 U3.

The **Transparent Snapshot Data Movers** pane provides a hierarchy view of the vCenter Server asset sources that have been added in PowerProtect Data Manager. Use this view to determine if the vCenter/ESXi is enabled for VIB management, and if the hosts have the VIB installed or are eligible for VIB installation. A vSphere host cluster can have one of the following statuses:

- **Installed**—The VIB installation on this vSphere host is completed, and TSDM is enabled as the default protection mechanism for the virtual machines on the vSphere host.
- **Ready for install**—The vSphere host requirements for VIB installation have been met, and the installation will proceed automatically on the vSphere host when a virtual machine running on the cluster is added to a protection policy.
- **Ready for upgrade**—This status displays when the VIB is installed on the vSphere host and PowerProtect Data Manager is upgraded, but the VIB is being managed manually. In this case, the VIB will not be upgraded automatically on the vSphere host.
- **Not eligible**—The vSphere host does not meet the requirements for VIB installation. When TSDM cannot be used, the VADP protection mechanism is used for virtual machine protection operations on this host.
- **Failed**—The VIB installation on the vSphere host did not complete successfully. The **Jobs** window provides more information about the issue that caused the failure.

Use the filter icon in the status column to display only vSphere hosts with a certain status. For example, you can choose to display only hosts that are ready for VIB installation or upgrade.

When the VIB installation is started, the **Protection Engines** window updates to display the progress. Also, an entry for the job **Performing Host Configuration (vib_install)** appears in the **Jobs** window.

- ① **NOTE:** Any virtual machine assets that were added to a virtual machine protection policy in PowerProtect Data Manager 19.8 and earlier currently use the VADP protection mechanism. After the VIB installation on the vSphere host that contains these virtual assets, you can migrate these assets to the TSDM protection mechanism. Migrating assets to use the Transparent Snapshot Data Mover on page 71 provides more information.

Disable or re-enable VIB on an ESXi host

In the PowerProtect Data Manager UI, you can disable VIB management on a vCenter to prevent automatic installation or update of the VIB on the ESXi host. To disable VIB management on the vCenter:

1. Go to **Infrastructure > Protection Engines**, and then select the **Transparent Snapshot Data Movers** pane.
2. Hover over the **Enabled** icon to the right of the vCenter, and then click **Disable**.

To re-enable VIB management on a vCenter that currently has the VIB disabled:

1. Hover over the icon to the right of the host, and then click **Enable**.

If a VIB installation or update is required, the status indicates **Ready for install** or **Ready for upgrade**.

2. Select the checkbox next to this host and click **Install** to manually perform the VIB install or update, or wait for the automatic VIB installation.
3. When performing a manual VIB installation, if one or more of the selections are not eligible or the VIB is already installed, a dialog appears. Click **OK** to proceed.

Migrating assets to use the Transparent Snapshot Data Mover

Transparent Snapshot Data Mover (TSDM) is the recommended protection mechanism for environments with vCenter/ESXi version 7.0 U3 and later installed, and is the default protection mechanism used for virtual machine assets protected by virtual machine crash-consistent policies in PowerProtect Data Manager 19.9 and later, provided that the policy is configured with the following options:

- **Performance optimization mode**.
- **Exclude swap files from backup** is off.
- **Enable guest file system quiescing** is off.

For existing virtual machine crash-consistent policies created with PowerProtect Data Manager version 19.8 and earlier, modifying the policy options to meet these requirements will migrate virtual machines on vSphere version 7.0 U3 and later clusters managed by a vCenter Server running version 7.0 U3 and later to use the TSDM protection mechanism.

You can also migrate virtual machine assets from the VADP protection mechanism to the TSDM protection mechanism by using the **Infrastructure > Assets** window of the PowerProtect Data Manager UI.

Before migrating assets to use TSDM, the vSphere Installation Bundle (VIB) is required. This installation occurs automatically, unless the use of TSDM is disabled on the vCenter Server asset source. Go to **Infrastructure > Protection Engines**, select the **Transparent Snapshot Data Movers** pane, and verify that the VIB is enabled on the vCenter. You can also expand the vCenter hierarchy view to confirm that the VIB installation has occurred on the vSphere hosts. Transparent Snapshot Data Mover protection mechanism on page 70 provides more information.

Migrate asset protection mechanism from VADP to TSDM

To migrate VADP virtual machine assets to use TSDM in the PowerProtect Data Manager UI:

1. Go to **Infrastructure > Assets** and select the **Virtual Machine** tab.
2. Filter the view to display the **Protection Mechanism** column.
3. Select one or more virtual machine assets with the VADP protection mechanism.
4. Select **More Actions > Protection Mechanism > Migrate to TSDM**.

Migrating assets to use the TSDM protection mechanism forces a new, full backup of these assets. This backup may take several minutes.

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 on page 60 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://sngosge.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**.

Managing Protection Policies

Topics:

- Protection policies
- Before you create a protection policy
- Supported enhanced VMware topologies for virtual-machine protection
- Add a protection policy for virtual-machine protection
- Add a Cloud Tier schedule to a protection policy
- Manual backups of protected assets
- Manual replication of protected assets
- Manual Cloud Tiering of protected assets
- Editing a protection policy
- View assets assigned to a protection policy
- Extended retention
- Edit the retention period for backup copies
- Delete backup copies
- Removing expired backup copies
- Removing assets from PowerProtect Data Manager
- Export protection
- Disable a protection policy
- Delete a protection policy
- Add a Service Level Agreement
- Export Asset Compliance
- Protection rules

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 and SQL databases
- Oracle databases
- SAP HANA databases
- File systems
- Kubernetes clusters
- Storage groups
- Network-attached storage (NAS)

This guide provides steps only for virtual-machine protection policies. For other policy types, refer to the individual user guides.

Before you create a protection policy

Consider the following best practices before creating a protection policy.

- An asset can be protected by only one policy at a time. Assets can be moved from one policy to another policy based on the priority of protection rules. In cases where protection rules result in assets moving from one policy to another, any assets that were manually selected for inclusion in the policy, however, will not be moved to a different policy.
- **NOTE:** If a SQL Server is hosted on a virtual machine, you can protect the SQL database with an application-consistent backup without interfering with the SQL 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 to avoid potential replication failures.
- Before adding replication to a protection policy, ensure that you add remote protection storage as the replication location. Add protection storage on page 39 provides detailed instructions about adding remote protection storage.
- Before you perform any backups on a weekly or monthly schedule from the protection policy, ensure that the PowerProtect Data Manager time zone is set to the local time zone.

Understanding backup terminology and managing backup frequency

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

- Different backup policy types can use different terminology to describe available backup levels. This terminology can differ not only between 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.

Refer to the following table to understand the different backup levels provided by each protection policy and to manage backup frequencies.

Table 26. Backup terminology and frequency

Protection-policy backup types	Available backup levels	Description	Equivalent traditional terminology	Minimum frequency recommendation
VMware application-aware	Full	Backs up all the blocks.	Full	Monthly
	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 actually 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
	Log	Backs up the transaction logs.	-	30 minutes
VMware crash-consistent	Full	Backs up all the blocks.	Full	Monthly
	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 actually 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

Table 26. Backup terminology and frequency (continued)

Protection-policy backup types	Available backup levels	Description	Equivalent traditional terminology	Minimum frequency recommendation
Kubernetes crash-consistent	Full	Backs up the namespace metadata and persistent volumes.	Full	Daily
	Synthetic Full	Backs up the namespace metadata, the blocks that have changed for persistent volumes on VMware first-class disks since the last synthetic-full or full backup, and all other persistent volumes in full. Although not all data has actually been 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	Backs up all the data.	Full	Monthly
	Synthetic Full	Backs up only the data that has 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 actually 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
Exchange centralized	Full	Backs up all the data.	Full	Weekly
	Synthetic Full	Backs up only the data that has 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 actually 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
SQL centralized	Full	Backs up all the data.	Full	Daily
	Differential	Backs up only the data that has changed since the last differential backup, or the last full backup if there are no other differential backups.	A differential backup is performed, followed by a merge operation that produces a full backup in storage.	12 hours
	Log	Backs up the transaction logs.	-	30 minutes
Oracle centralized	Full	Backs up all the data.	Full	Daily
	Incremental Cumulative	Backs up only the data that has changed since the last full backup.	Differential	12 hours

Table 26. Backup terminology and frequency (continued)

Protection-policy backup types	Available backup levels	Description	Equivalent traditional terminology	Minimum frequency recommendation
	Incremental Differential	Backs up 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.	Incremental	6 hours
	Log	Backs up the archived logs.	-	30 minutes
SAP HANA centralized	Full	Backs up all the data.	Full	Daily
	Differential	Backs up only the data that has changed since the last full backup.	Differential	12 hours
	Incremental	Backs up only the data that has changed since the last incremental backup, or the last full backup if there are no other incremental backups.	Incremental	6 hours
VMAX storage group centralized	Full	Backs up all the blocks.	Full	Daily
	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 actually 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 this include, but are not limited to, 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.

Supported enhanced VMware topologies for virtual-machine protection

PowerProtect Data Manager provides protection for clustered ESXi server storage, networking, and enterprise management. Understanding what topologies are supported in these environments aids in the design of your network infrastructure.

Supported enhanced topologies

Supported topologies of clustered ESXi server storage, networking, and enterprise management include the following:

- vSAN operations
- NSX-T port groups

- Enhanced Link Mode vCenter servers

For more information see the E-Lab Navigator:

vSAN operations

Standard clusters, stretched clusters, two-node clusters, and HCI Mesh datastores support the following operations:

- Backing up and restoring virtual machines
- Search Engines
- VM Direct Engines
- HA failover of Search Engines and VM Direct Engines
- Post-failover protection

NSX-T port groups

PowerProtect Data Manager supports the use of NSX-T with up to 2,000 port groups. These can be default VDS port groups or N-VDS port groups, and they support the following components:

- PowerProtect Data Manager servers
- VM Direct Engines
- Search nodes
- Workload virtual machines

Enhanced Link Mode vCenter servers

Enhanced Linked Mode connects multiple vCenter Server systems together by using one or more Platform Services Controllers (PSCs). PowerProtect Data Manager supports the protection of workload virtual machines running inside Enhanced Linked Mode vCenter servers. This protection also applies during and after any vMotion operation of the virtual machines.

To support virtual machine protection workflows for vCenter Servers that are in Enhanced Linked Mode, PowerProtect Data Manager requires you to add all of the linked vCenters as asset sources, and also to install the PowerProtect **vSphere Plugin** on all of these vCenters

Add a protection policy for virtual-machine protection

A protection policy enables you to select a specific group of assets that you want to back up and replicate. Perform the following steps to create a virtual-machine protection policy in the PowerProtect Data Manager UI.

Prerequisites

It is recommended that you distribute virtual-machine asset protection workloads over multiple ESXi hosts so that you do not exceed the ESXi Network Block Device (NBD) session limit. If the limit is reached, you can manage the workload by deploying an external VM Direct Engine on the host or cluster using **Hot Add** transport mode. Also, Dell Technologies recommends during policy configuration to assign virtual machines to a protection policy based on logical grouping to allow for better scheduling of backups. Grouping helps avoid resource contention and creates more organized logs for review.

To create application-aware protection policies for virtual machines, ensure that:

- You manually update the VMX configuration parameter `disk.EnableUUID` to `True` by using the **vSphere Web Client**.
- The vSphere version that you are running uses a supported version of VMware Tools. Software compatibility information for the PowerProtect Data Manager software is provided in the e-Lab Navigator, available at <https://elabnavigator.emc.com/eln/modernHomeDataProtection>.
- The virtual machine has direct access to the DD client.
- The virtual machine uses SCSI disks only, and the number of available SCSI slots matches at least the number of disks.
- The Windows account that is used for the protection policy is limited to the local system Administrator or the domain Administrator. This user requires both Microsoft Windows administrative rights and Microsoft SQL Server login and sysadmin rights.

- SQL configuration support is limited to Microsoft SQL Server stand-alone instances and a Microsoft SQL Server Always On availability group (AAG) configured with file share witness. Unsupported configurations include Microsoft SQL Server failover cluster instances that are configured with shared drives, and Microsoft SQL Server cluster-less AAG configurations.
- For Microsoft SQL Server AAG configurations, the database administrator specifies the AAG backup preferences for backup in the Microsoft SQL Server Management Studio (SSMS). These preferences control which AAG node is selected as the preferred node when you perform a transaction log backup of AAG databases.
- To protect virtual machines that use virtualization-based security (VBS) and virtual Trusted Platform Module 2.0 (vTPM), vCenter 7.0 U1 or later is required.

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

Chapter Managing Storage on page 38 provides more information about working with storage units, including applicable limitations and security considerations.

Before performing any backups on a weekly or monthly schedule from the protection policy, ensure that the PowerProtect Data Manager time zone is set to the local time zone.

About this task

For virtual-machine protection policies, data is moved using one of two types of protection mechanisms:

- **Transparent Snapshot Data Mover**—Starting in PowerProtect Data Manager version 19.9, Transparent Snapshot Data Mover (TSDM) is the default protection mechanism that is used for crash-consistent virtual-machine policies when the following requirements are met:
 - vCenter/ESXi version 7.0 U3 and later is installed in the environment.
 - The protection policy uses **Performance** optimization mode, with the **Exclude swap files from backup** and **Enable guest file system quiescing** checkboxes cleared.
- **VADP**—VMware vStorage API for Data Protection (VADP) is the protection mechanism that is used for application aware virtual-machine policies and crash-consistent policies that do not meet the TSDM software requirements. VADP is the only protection mechanism available in PowerProtect Data Manager versions 19.8 and earlier.

The section Transparent Snapshot Data Mover protection mechanism on page 70 provides more information about TSDM.

Steps

1. From the left navigation pane, select **Protection > Protection Policies**.
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 **Virtual Machine**, which includes protection for SQL application-aware virtual machines.
4. On the **Purpose** page, select from the following options to indicate the purpose of the new protection policy group, and then click **Next**:
 - **Crash Consistent**—Select this type for point-in-time backup of virtual machines.
 - **Application Aware**—For virtual machines with a SQL application installed, select this type to quiesce the application to perform the SQL database and transaction log backup. When you select this type, you must provide Windows account credentials for the virtual machine. You can provide the credentials at the protection-policy level or the virtual-machine asset level. When you provide the credentials at both levels, the virtual-machine asset credentials override the policy credentials.
 - **Exclusion**—Select this type if there are virtual-machine assets within the protection policy that you plan to exclude from data protection operations.

By default, quiescing is automatically performed for the guest file system on the virtual machine. Quiescing ensures that the data within the guest file system is in a state that is appropriate for backups. If the file system cannot be quiesced on the first attempt, then the snapshot and backup are performed without quiescing.

VMware Tools is used to quiesce the file system in the guest operating system. The VMware documentation provides more information.

5. On the **Assets** page, select the assets for inclusion in this policy by choosing one of the following options from the list:
 - **View by Host**—This option enables you to view all assets within a specific host, and then select individual assets or a group of assets at a host or container level for policy inclusion. For example:
 - Select a stand-alone host to include all assets under this host.

- ① **NOTE:** If you select a host in a cluster, no assets are selected. For a host in a cluster, ensure that you select the cluster or other containers (for example, a resource pool or vApp) under the cluster host.
- o Expand the tree and select a container level in the vCenter hierarchy (for example, the data center, cluster, host, or resource pool) to include all assets under that level. If assets at any level are protected by another policy, a label with the name of that policy appears next to the level.

① **NOTE:** VMs created by the vSphere Cluster Service (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 `vadm-discovery.log` provides a list of vCLS VMs that are excluded from protection.

When you select a container level in the **View by Host** view, a protection rule is automatically created to ensure that these container level selections will be retained, even if changes occur from movements within the vSphere environment or the names of resource pools or folders change. This rule is managed by the PowerProtect Data Manager system, and cannot be modified. The rule will also be updated automatically if you make changes to container selections when editing the policy, or when assets are moved into or out of a selected container.

To view this rule after policy creation, go to **Protection > Protection Rules**. The name in the **Protection Rule Name** column for this new rule matches the policy name.

If this new rule results in an overlap of protection with an existing rule, you can resolve these conflicts by changing the policy protection rule priority in the **Selection Overlap** page. Step 7 on page 79 provides more information.

- ① **NOTE:** The behavior of automatic rule creation that allows assets to move into or out of policies can only be modified in the REST API. After updating from a previous release, if **View by Host** is not visible you can enable this view by manually changing the `/api/v2/common-settings/DYNAMIC_FILTER_SETTING`. The PowerProtect Data Manager Public REST API documentation provides instructions.
- o Expand the tree and select individual assets within containers.

When you select individual assets within this view, these selections are considered static, and no protection rule is automatically created. In cases where protection rules result in assets moving from one policy to another, any assets that are manually selected for inclusion in the policy will not be moved to a different policy.

- **View Asset Table**—This option enables you to view all unprotected assets in the vCenter within a table, and then select individual unprotected assets that you want to back up as part of this protection policy. In cases where protection rules result in assets moving from one policy to another, any assets that are manually selected for inclusion in the policy will not be moved to a different policy.

When you select a virtual-machine asset in this view, a dialog displays indicating that you can exclude virtual disks (VMDKs) from protection of these assets. To dismiss the dialog for other selections, select the check box and click **OK**.

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. If the virtual machines that you want to protect are not listed, use the **Search** box to search by asset name.

- ① **NOTE:** When you configure a virtual-machine application-aware protection policy to protect a Microsoft SQL Server Always On availability group (AAG), you must add all the virtual machines for that AAG to the same policy, to ensure proper protection. Failure to do so might result in missed transaction log backups.

For the virtual-machine application-aware case, the **Assets** page displays a warning about the AAG policy configuration requirement.

6. Optionally, if you want to exclude nonproduction 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**.
 - ① **NOTE:** For PowerProtect Data Manager version 19.3, a virtual machine with disk exclusion and Cloud Disaster Recovery (DR) cannot co-exist in the same protection policy. If you exclude disks from a virtual-machine protection policy, Cloud DR is not supported.
 - 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**.

If any virtual objects or assets that were selected in the previous page overlap with assets that are already protected by another policy, the **Selection Overlap** page appears. Overlap can occur, for example, when two policies (the new policy and an existing policy) use the **View by Host** view for asset selection by container level.

- a. To switch protection of any virtual objects listed in the **Protection Priority Overlap** table from an existing policy, update the **Policy Priority** field to a level equal to or higher than the other policy currently protecting these objects. The lower the value, the higher the priority. For example, **1** is the highest priority. When you change this value, the priority of the rule that is associated with this policy is also changed.
- b. To switch protection of any assets that are listed in the **Asset Protection Overlap** table to this policy, select the checkbox next to one or more assets. Note that selecting these assets for inclusion in this policy removes the assets from the other policy.

When you change the priority of the selected assets, the protection rule is updated automatically.

8. Click **Next**.
The **Objectives** page appears.
9. On the **Objectives** page, 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.
Add a Service Level Agreement on page 102 provides instructions.

10. Click **Add** under **Primary Backup**.
The **Add Primary Backup** dialog appears.

11. On the **Schedules** pane of the **Add Primary Backup** dialog:
 - a. Specify the following fields to schedule the synthetic full backup of this protection policy:
 - **Create a Synthetic Full...**—Specify how often to create a synthetic full backup. A **Synthetic Full** backs up only the changed blocks since the last backup to create a new full backup.
 - **Retain For**—Specify the retention period for the synthetic full backup.

You can extend the retention period for the latest primary backup copy by using the **Extend Retention** schedule. For example, your regular schedule for daily backups can use a retention period of 30 days, but you can apply extended retention to keep the full backups taken on Mondays for 10 weeks. Step 14 on page 81 provides instructions.

i **NOTE:** For database backups, PowerProtect Data Manager chains the dependent backups together. For example, the synthetic full or transaction log backups are chained to their base full backup. The backups do not expire until the last backup in the chain expires. This ensures that all synthetic full and transaction log backups are recoverable until they have all expired.

- **Start and End**—For 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.

i **NOTE:** Any backups started before the **End Time** occurs continue until completion.

- Click **Save** to save and collapse the backup schedule.

- b. Click **Add Backup** if you want to periodically force a full (level 0) backup, and then specify the following fields to schedule the full backup of this protection policy:

i **NOTE:** When you select this option, the backup chain is reset.

- **Create a Full...**—Specify whether you want to create a weekly or monthly full backup.
- **Repeat on**—Depending on the frequency of the full backup schedule, specify the day of the week or the date of the month for the full backup.
- **Retain For**—Specify the retention period for the full backup. This can be the same value as the synthetic full backup schedule, or a different value.
- **Start and End**—For the activity window, specify a time of day to start the full backup, and a time of day after which backups cannot be started.

i **NOTE:** Any backups started before the **End Time** occurs continue until completion.

- Click **Save** to save and collapse the backup schedule.

- c. For virtual-machine application-aware protection policies, click **Add Backup** to create a log backup, and then specify the following fields:

- **Create a Log...**—For application-aware protection policies, specify the interval in minutes for log generation.

i **NOTE:** For SQL Server AAG configurations, the database administrator can specify the AAG backup preferences for a transaction log backup in the Microsoft SQL Server Management Studio.

- **Retain For**—Specify the retention period for the log backup. This can be the same retention value specified for the synthetic full or full schedule, or a different value.

i **NOTE:** Setting a shorter retention period for log backups than the full backup can result in data loss and the inability to restore point-in-time copies.

- **Start and End**—For the activity window, specify a time of day to start the log backup, and a time of day after which log backups cannot be started.

i **NOTE:** Any backups started before the **End Time** occurs continue until completion.

- Click **Save** to save and collapse the backup schedule.

12. On the **Target** pane of the **Add Primary Backup** dialog, specify the following fields:

- a. **Storage Name**—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 **Storage Target** window.

i **NOTE:** The **Space** field indicates the total amount of space, and the percentage of available space, on the protection storage system.

- 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, **testvmplc-ppdm-daily-123ab (300 GB/1 TB)**

When you select **New**, a new storage unit in the format *policy name host name unique identifier* is created in the storage system upon policy completion. For example, **testvmplc-ppdm-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 on the selected system. PowerProtect Data Manager uses Governance mode for retention locking, which means that the lock can be reverted at any time if necessary. Moving the **Retention Lock** slider on or off applies to the current backup copy only, and does not impact the retention lock setting for existing backup copies.

i **NOTE:** Primary backups are assigned a default retention lock period of 14 days. Replicated backups, however, are not assigned a default retention lock period. If you enable **Retention Lock** for a replicated backup, ensure that you set the **Retain For** field in the **Add Replication** backup schedule dialog to a minimum number of 14 days so that the replicated backup does not expire before the primary backup.

- e. **SLA**—Select an existing service level agreement that you want to apply to this schedule from the list, or select **Add** to create an SLA within the **Add Service Level Agreement** wizard.

Add a Service Level Agreement on page 102 provides instructions.

13. Click **Save** to save your changes and return to the **Objectives** page.

The **Objectives** page updates to display the name and location of the target storage system under **Primary Backup**.

i **NOTE:** After completing a backup schedule, you can change any schedule details by clicking **Edit** next to the schedule.

14. Optionally, extend the retention period for the latest primary backup copy:

Extended retention on page 93 provides more information about **Extend Retention** functionality.

- a. Click **Extend Retention** next to **Primary Backup**. An entry for **Extend Retention** is created below **Primary Backup**.

- b. Under **Extend Retention**, click **Add**. The **Add Extended Retention** dialog appears.

- c. **Retain the next scheduled full copy every...**—Specify a weekly, monthly, or yearly recurrence for the extended retention backup schedule.

- d. **Repeat on**—Depending on the frequency of the full backup schedule, specify the day of the week, the date of the month, or the date of the year that the extended retention backup will occur.

- e. **Retain For**—Specify the retention period for the backup. You can retain an extended retention backup for a maximum of 70 years.

- f. Click **Save** to save your changes and return to the **Objectives** page.

15. Optionally, replicate the full and synthetic full backups to a remote storage system:

- a. Click **Replicate** next to **Primary Backup** or **Extend Retention**. An entry for **Replicate** is created to the right of the primary or extended retention backup schedule.

i **NOTE:** PowerProtect Data Manager supports replicating an extended retention backup only if the primary backup already has one or more replication stages. Also, for replication of an extended retention backup, you can only select the protection storage systems that are used by the replication stages based on the primary stage.

For example, if there are 6 systems available (DD001-DD006), and the primary backup is on DD0001:

- Replicate1 based on the primary backup is replicated to DD002
- Replicate2 based on the primary backup is replicated to DD003

- Extended retention backup is backed up to DD001
- Replicate3 based on the extended retention backup must be replicated to DD002 or DD003.

b. Under **Replicate**, click **Add**. The **Add Replication** dialog appears.

i **NOTE:** To enable replication, ensure that you add remote protection storage as the replication location. Add protection storage on page 39 provides detailed instructions about adding remote protection storage.

c. Complete the schedule details in the **Add Replication** dialog, and then click **Save** to save your changes and return to the **Objectives** page.

The schedule frequency can be every day, week, month, or x hours for replication of the primary backup, and every day, week, month, year, or x hours for replication of the extended retention backup. 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 1 to 12 hours.

All replication copies of the primary backup schedule will use the same retention period, and by default, this retention period is inherited from the **Retain For** value of the synthetic full backup schedule. To specify a different retention period for all of the replication copies of this primary backup schedule, click **Edit**, change the value in the **Retain For** field, and then click **Save**. This retention period will be applied to all of the replicated copies (synthetic full and full) of this primary backup schedule.

When creating multiple replication copies of the same protection policy, Dell Technologies recommends selecting a different storage system for each copy.

16. Optionally, to move backups from DD storage to Cloud Tier, add a Cloud stage for the primary, replication, or extended retention schedule:

a. Click **Cloud Tier** next to **Primary Backup** or **Extend Retention** or, if adding a Cloud stage for a replication schedule that you have added, click **Cloud Tier** under **Replicate**. An entry for **Cloud Tier** is created to the right of the primary or extended retention backup schedule, or below the replication schedule.

b. Under the entry for **Cloud Tier**, click **Add**.

The **Add Cloud Tier Backup** dialog appears, with summary schedule information for the parent node to indicate whether you are adding this Cloud Tier stage for the primary backup schedule, the extended retention backup schedule, or the replication schedule.

c. Complete the schedule details in the **Add Cloud Tier Backup** dialog, and then click **Save** to save your changes and return to the **Objectives** page.

Add a Cloud Tier schedule to a protection policy on page 86 provides detailed instructions for adding a Cloud stage for a primary, replication, or extended retention schedule.

i **NOTE:** In order to move a backup or replica to Cloud Tier, schedules must have a retention time of 14 days or more. Also, discovery of protection storage that is configured with a Cloud unit is required.

17. Optionally, if **Cloud Disaster Recovery** is configured in the **Infrastructure > Storage** window, you can add a Cloud DR stage for virtual-machine protection policies:

a. Click **Cloud DR** next to **Primary Backup** or **Extend Retention** or, if adding a Cloud stage for a replication schedule that you have added, click **Cloud DR** under **Replicate**. An entry for **Cloud DR** is created to the right of the primary or extended retention backup schedule, or below the replication schedule.

b. Under the entry for **Cloud DR**, click **Add**.

The **Add Cloud DR Backup** dialog appears, with summary schedule information for the parent node to indicate whether you are adding this Cloud DR stage for the primary backup schedule, the extended retention backup schedule, or the replication schedule.

c. Complete the schedule details in the **Add Cloud DR Backup** dialog, and then click **Save** to save your changes and return to the **Objectives** page.

The *PowerProtect Data Manager Cloud Disaster Recovery Administration and User Guide* provides detailed instructions for adding a Cloud DR stage for a primary, replication, or extended retention schedule.

18. Click **Next**.

The **Options** page appears.

19. On the **Options** page:

a. For **Optimize For**, select from one of the following backup optimization modes:

- **Performance**—Optimize for backup and replication speed. Selecting this mode results in more storage consumption. Previous versions of PowerProtect Data Manager used this option by default.

When using the **Transparent Snapshot Data Mover** protection mechanism, select **Performance** optimization mode.

- **Capacity**—Optimize for backup size. Selecting this mode results in less storage consumption, but backups take longer to complete.

NOTE: Changing the optimization mode after the first backup of the protection policy forces the next backup to be a full backup, and results in increased storage capacity usage due to differences in how each mode uses data deduplication. This increase continues until all backups performed using the previous optimization mode expire and have been deleted.

- b. **Exclude swap files from backup**—Select to exclude the C:\swapfile.sys, C:\pagefile.sys, and C:\hiberfil.sys swap and memory files of Microsoft Windows virtual machines, in the virtual-machine backup. By default, this checkbox is cleared.

When using the **Transparent Snapshot Data Mover** protection mechanism, do not select the **Exclude swap files from backup** checkbox.

NOTE: Including swap and memory files in a backup unnecessarily increases the size of the backup and the time to RTO during recovery. These files are rebuilt by the Microsoft Windows operating system upon restart, and not required for recovery.

- c. **Enable indexing for file search and restore**—Select to enable indexing. This option is visible only upon activating the search cluster node.
- d. **Enable guest file system quiescing**—Select to enable **VMware Tools** to quiesce the file system during crash-consistent virtual-machine backups.

When using the **Transparent Snapshot Data Mover** protection mechanism, do not select the **Enable guest file system quiescing** checkbox.

20. Click **Next**.

The **Summary** page appears.

21. Review the protection policy group configuration details. Except for the protection policy type, you can click **Edit** next to any details to change the protection policy information. When satisfied with the details, click **Finish**.

An informational message appears to confirm that PowerProtect Data Manager has saved the protection policy.

When the new protection policy is created and assets are added to the protection policy, PowerProtect Data Manager performs backups according to the backup schedule.

For virtual machines, if you have not yet added a VM Direct Engine, the backup is performed using the embedded VM Direct Engine that is included with PowerProtect Data Manager. Subsequent backups are performed according to the schedule specified.

NOTE: If the target virtual-machine datastore for backup is running low on free space and the datastore free space threshold is configured in **vCenter Settings**, a warning message appears or a backup failure occurs. When the **Datastore Free Space Warning Threshold** is reached, the backup proceeds with a warning message in the logs. When the **Datastore Free Space Failure Threshold** is reached, the backup fails.

To check the warning and failure threshold values, select **Infrastructure > Asset Sources** and click the **vCenter** tab. Click the gear icon to open the **vCenter Settings** dialog.

22. 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.

Managing virtual-machine backups

The following sections describe the options that are available for virtual-machine assets that are backed up as part of a protection policy.

Add and remove the credentials for virtual-machine assets

You can optionally add and remove the credentials for multiple virtual-machine assets at the same time in the PowerProtect Data Manager UI. With previous versions, you could add and remove the credentials for one virtual-machine asset at a time.

About this task

NOTE: The asset-level credentials take precedence over policy-level credentials for virtual machines. Asset-level credentials have the highest precedence. Virtual machines do not support the asset source-level (host) credentials.

Use the following procedure to add or remove one or more credentials for virtual-machine assets.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Assets**, and then click the **Virtual Machine** tab.
2. Select one or more assets by clicking the checkbox next to each required asset name.
3. Select **More Actions > Set Credential**.
4. In the **Set Credential** dialog box, add or remove the credentials for the selected virtual-machine assets:
 - To add the credentials for the assets, select the appropriate value from the drop-down list in the **Credential** field:
 - To create new credentials, select **Create New**.
In the **Add Credentials** dialog box that appears, specify the required field values and then click **Save**.
 - To add existing credentials, select the credentials name from the credentials list.
 - To remove the credentials for the assets, select **Remove Credentials**.
5. Click **Save** in the **Set Credential** dialog box.

Results

After you add the credentials by using this procedure, the asset-level credentials are used for the selected assets during the virtual-machine centralized backups, overriding the policy-level credentials.

Enable or disable Changed Block Tracking (CBT)

The Changed Block Tracking (CBT) feature is used to identify areas of the virtual-machine backup that have changed since the last backup and only process those changed areas during the next backup. CBT is enabled by default.

About this task

To set Changed Block Tracking (CBT) for virtual machines, complete the following steps:

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
2. From the **Assets** window, select the **Virtual Machine** tab. If a policy has been assigned, the table lists the virtual-machine assets that have been discovered in the vCenter, along with the associated protection policy.
3. Select one or more virtual-machine assets from the list, and click **More Actions > Changed Block Tracking**.
The **Changed Block Tracking** dialog box appears.
4. Clear the check box to disable CBT, or select the check box to enable CBT.
In some cases, CBT can cause backups to take longer than expected if there are high change rates on the virtual machine. You can disable CBT for virtual machines if the backups are taking too long to complete. Also, if you encounter an issue with CBT, you can disable it on the virtual machine.

NOTE: If CBT is enabled in PowerProtect Data Manager but is disabled in VMware vSphere, PowerProtect Data Manager tries to back up the virtual machine with CBT enabled. If PowerProtect Data Manager cannot enable CBT, the backup completes with a warning that indicates CBT data is not available.

5. Click **Save**.

NOTE: When CBT is disabled for a virtual machine, subsequent backups no longer use CBT.

More options for managing virtual-machine backups

After you create a virtual-machine protection policy, additional options become available for virtual-machine assets that are backed up as part of the policy.

To access these options:

1. From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
2. From the **Assets** window, select the **Virtual Machine** tab. If a policy has been assigned, the table lists the virtual-machine assets that have been discovered in the vCenter, along with the associated protection policy.

NOTE: 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**.

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 the storage icon to the right of the VM icon, for example, **DD**. The table in the right pane lists the backup copies.

Depending on whether the asset is retention locked, you can perform the following functions from this window:

- Edit the retention period of backup copies to extend or shorten the amount of time that backups are retained—Select one or more backup copies from the table and click **Edit Retention**.
 - To select a calendar date as the expiration date for backups, select **Retention Date**.
 - To define a fixed retention period in days, weeks, or months 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.

- Delete a backup copy—If you no longer require a copy and the retention lock is not enabled, select the copy from the table and click **Delete**.

Snapshot freeze scripts and thaw scripts for virtual-machine backups

You can use custom scripts to back up a Windows or Linux virtual machine which runs an application that PowerProtect Data Manager does not directly support. These scripts run before and after the snapshot to place the virtual machine and application into a state where you can perform a backup.

NOTE: Use of these scripts is not supported for virtual machines with the Transparent Snapshot Data Mover (TSDM) protection mechanism enabled.

Table 27. Script descriptions and related terms

Script	Related terms		Description
Freeze	Quiesce	Pre-freeze	This script runs before the snapshot initialization to quiesce the virtual machine and place the application in a frozen state. Quiescing ensures that the data within the guest file system is in a consistent state that is appropriate for backups.
Thaw	Unquiesce	Post-thaw	This script runs after the snapshot finalization to unquiesce the virtual machine, thaw the application, and then return the virtual machine to normal operation.

PowerProtect Data Manager uses the VMware Tools package to quiesce the virtual machine. The VMware documentation provides more information. Before you deploy the freeze and thaw scripts, install the latest version of the VMware Tools package on the virtual machine.

The freeze and thaw scripts are specific to each application. If the freeze script returns a nonzero exit code, snapshot creation fails.

After you create your custom scripts, deploy the scripts to the correct location on the virtual machine, as specified in the following tables.

Table 28. Script locations for Windows virtual machines

ESXi version	Freeze script location	Thaw script location
ESXi 6.5 or later	C:\Program Files\VMware\VMware Tools\backupScripts.d\ All scripts are invoked in ascending alphabetical order with <code>freeze</code> as the first argument.	C:\Program Files\VMware\VMware Tools\backupScripts.d\ All scripts are invoked in descending alphabetical order with <code>thaw</code> or <code>freezeFail</code> as the first argument.

Table 29. Script locations for Linux virtual machines

ESXi version	Freeze script location	Thaw script location
ESXi 6.5 or later	<code>/usr/sbin/pre-freeze-script</code>	<code>/usr/sbin/post-thaw-script</code>

For Linux virtual machines, set the script ownership and permissions after you deploy the scripts:

- `sudo chown root:root /usr/sbin/pre-freeze-script /usr/sbin/post-thaw-script`
- `sudo chmod 0700 /usr/sbin/pre-freeze-script /usr/sbin/post-thaw-script`

Add a Cloud Tier schedule to a protection policy

For some protection policy types, you can add a Cloud Tier schedule to a protection policy in order to perform backups to Cloud Tier.

Prerequisites

Ensure that a protection storage system is set up for Cloud Tiering.

About this task

You can create the Cloud Tier schedule from **Primary Backup**, **Replicate**, and **Extend Retention** stages. Schedules 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.

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

- Kubernetes
- On the **Purpose** page, select from the available options to indicate the purpose of the new protection policy, and then click **Next**.
 - On the **Assets** page, select the assets that you want to protect with this policy, and then click **Next**.
 - On the **Objectives** page, click **Add** under **Primary Backup** if the primary backup schedule 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 stage, however, the Cloud Tier schedule requires a minimum retention period of 14 days in the **Retain for** field.
 - Click **Cloud Tier** next to **Primary Backup** or **Extend Retention** or, if adding a Cloud stage for a replication schedule that you have added, click **Cloud Tier** under **Replicate**.
An entry for **Cloud Tier** is created to the right of the primary backup or extended retention schedule, or below the replication schedule.
 - Under the entry for **Cloud Tier**, click **Add**.
The **Add Cloud Tier Backup** dialog appears, with summary schedule information for the parent node. This information indicates whether you are adding this Cloud Tier stage for the primary backup schedule, the extended retention schedule, or the replication schedule.
 - In the **Add Cloud Tier Backup** dialog box, set the following parameters and then click **Save**:
 - Select the appropriate storage unit from the **Cloud Target** list.
 - For **Tier After**, set a time of 14 days or more.

The protection policy schedule is now enabled with Cloud Tiering.
 - 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

- From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
- Select an asset and click **View Copies**.
- Click an asset copy icon.
Cloud Tier backups are listed by cloud storage in the **Location** column.
- To change how long copies remain in cloud storage, complete the following steps:
 - Select a Cloud Tier backup and click **Edit Retention**.
 - 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 8 months.
 - 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.
- 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 on page 95 and Remove backup copies from the PowerProtect Data Manager database on page 97 provides more information.
- 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.
- To extend the date to re-tier the copy back to the cloud, select **Edit Recall Retention**.

- To manually move a copy back to cloud storage, select **Retier**.

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.

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. To perform this manual backup:

- From the PowerProtect Data Manager UI, select **Protection > Protection Policies**
- Select the protection policy that contains the assets that you want to back up, and click **Protect Now**.

The **Protect Now** wizard appears.

NOTE: The protection policy must be enabled, and its purpose must not be Exclusion or Self-Service Protection.

- 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. 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.

- On the **Configuration** page, select **Back up now**, and then select from the available backup types.
- 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 stage of the parent protection policy.

- On the **Summary** page, review the settings and then click **Protect Now**. A notification appears indicating whether the request was processed successfully.

You can also perform a manual backup from the **Infrastructure > Assets** window, but only for one asset at a time. To perform this manual backup:

- From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
- Select the tab for the asset type you want to back up. A list of assets appears.
- Select an asset from the table that has an associated protection policy.

NOTE: You can select only one asset at a time for manual backup. The protection policy must be enabled, and its purpose must not be Exclusion or Self-Service Protection. A full backup is created for the selected asset.

- Click **Protect Now**. A notification appears indicating whether the request was processed successfully.

When a virtual machine is part of an application-aware protection policy, the manual backup is a full application-aware backup.

NOTE: The manual backup is managed by other configured stages (extended retention backup, replication, Cloud Tier, Cloud DR) of the parent protection policy.

Manual replication of protected assets

You can perform replication of one more protected assets within a protection policy by using the **Protect Now** functionality in the PowerProtect Data Manager UI.

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.

To perform manual replication:

- From the PowerProtect Data Manager UI, select **Protection > Protection Policies**
- Select the protection policy that contains the assets that you want to replicate, and click **Protect Now**.

The **Protect Now** wizard appears.

NOTE: The protection policy must be enabled, its purpose must not be Exclusion, and the policy must already be configured with a replication stage.

- On the **Assets Selection** page, select whether you want to replicate 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 replication instead of all assets, the **Assets** page appears with the individual assets available for selection. Select the assets that you want to include in the manual replication, and then click **Next** to display the **Configuration** page.

If you selected to replicate all assets, the **Configuration** page appears.

- On the **Configuration** page, select **Replicate now**, and then select from the available replication stages.

NOTE: Only replication stages for the primary backup are available for selection.

- 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 stage of the parent protection policy.

- 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 stage, 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 stage.

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.
- Click the **DD** icon 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.

Editing a protection policy

You can use the PowerProtect Data Manager UI to change any of the following information for an existing enabled or disabled protection policy:

- Policy name and description
- Backup schedule
- Backup optimization mode
- Settings for network interface, storage quotas, and retention lock
- Adding or removing assets from the policy.

You cannot modify a protection policy type or purpose. For these actions, add 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.

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](#) on page 91.

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.
NOTE: Dell Technologies recommends that you do not edit the network interface for application agent assets such as File System, SQL, ORACLE, and SAP HANA, because modifying this setting causes subsequent backup failure. As a workaround, set the lockbox, which initiates a new asset configuration.

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 and storage units](#) on page 90.
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 and storage units

You can change the storage target or storage unit that PowerProtect Data Manager targets for each protection policy.

When editing protection policies in the PowerProtect Data Manager UI, for the **Primary Backup** and **Replicate** schedules on the **Objectives** page:

- The **Storage Name** drop-down list shows the current storage target and those storage targets that are available for the protection policy.
- The **New** and **Existing** options for **Storage Unit** show the current storage unit that PowerProtect Data Manager targets on the selected protection storage system.

Storage targets

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

- The selected storage target for **Storage Name** in the **Primary Backup** row does not appear in the drop-down list for **Storage Name** in the **Replicate** row.
- The selected storage target for **Storage Name** in the **Replicate** row does not appear in the drop-down list for **Storage Name** in the **Primary Backup** row.

- Only those storage targets that have been licensed and configured for use by the current protection policy appear in a drop-down list.
- If a storage target exists but does not appear in a drop-down list, click **Add** at the bottom of the list. Configure the storage target for use with the protection policy.
- When a storage target is changed, a new `storage_unit_name` is automatically created, and the configuration is passed to any backup agents.
- Changing storage targets in Storage Group protection policies is not supported.

NOTE:

Changing the storage target for **Primary Backup** may prevent any scheduled backups from being performed until after the next full backup. To ensure that all scheduled backups are performed on schedule, click **Back Up Now** from the **Protection > Protection Policies** pane.

This guidance does not apply to VMware crash-consistent or file system backups for **Primary Backup**. For those asset types, you can change the storage target and all scheduled backups happen without further action. This guidance also does not apply to replication.

When you change a storage target, appropriately configure any dependencies. For example, configuring a cloud provider to use the new storage target.

Storage units

Storage units on page 40 provides more information about working with storage units, including applicable limitations and maintenance considerations.

When you select **New**, PowerProtect Data Manager maintains a dedicated storage unit for this protection policy. Click **Set Storage Quotas**.

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.

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.

- **Capacity Quota**—Controls the total size of precompression data that is written to the protection storage.
- **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.

When you select **Existing**, the protection policy targets a storage unit under the control of PowerProtect Data Manager. Click **Select**.

The **Select Storage Unit** dialog box opens and displays a list of the storage units under the control of PowerProtect Data Manager.

Select a storage unit from the list, and then click **Select**.

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.
 - ① **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.
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**.
 - ① **NOTE:** For PowerProtect Data Manager version 19.3, a virtual machine with disk exclusion and Cloud Disaster Recovery (DR) cannot coexist in the same protection policy. If you exclude disks from a virtual machine protection policy, Cloud DR is not supported.
 - 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.
9. 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.

View assets assigned to a protection policy

You can view assets that are assigned to a protection policy. If the modification of a protection rule results in assets moving 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 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.

Extended retention

You can extend the retention period for the primary backup copy for long term retention. For example, your regular schedule for daily backups can use a retention period of 30 days, but you can extend the retention period to keep the 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. When you extend the retention period of a backup in a protection policy, you 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 stage for a protection policy, you define a set of 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 an ad-hoc 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 schedule configuration.

When you configure extended retention to match backups by weekday, PowerProtect Data Manager may identify a backup that was taken on one weekday as being taken on a different 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 that 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 that 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 that 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, automatic extended retention creates a job at the beginning of the backup window for the primary stage. This job remains queued until the end of the backup window and then starts.

The following examples describe the behavior of backups with extended retention for centralized and self-service protection.

Centralized protection

For an hourly primary backup schedule that starts on Sunday at 8:00 p.m. and ends on Monday at 6:00 p.m. with a weekly extended retention schedule that is 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. For a backup schedule that starts on Sunday at 12:00 p.m. and ends on Monday at 12:00 p.m. with a weekly extended retention schedule that is 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.

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.

Steps

1. From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
2. 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.

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**.

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 the storage icon to the right of the icon for the asset, for example, **DD**. The table in the right pane lists the backup copies.
 5. Select one or more backup copies from the table and click **Edit Retention**.
 6. 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.
7. 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.

Delete backup copies

In addition to deleting backups upon expiration of the retention period, 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

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 the storage icon to the right of the icon for the asset, for example, **DD**. The table in the right pane lists the backup copies.
 5. 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 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.
6. 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**.
 7. 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**.

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 is 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 is shown in UTC. Go to **Alerts > Audit Logs** to view the audit log.

8. 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

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 the storage icon to the right of the icon for the asset, for example, **DD**. 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 **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 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 successfully deleted from protection storage. If the deletion is successful, the deleted copies no longer appear in the table.

Export data for deleted 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 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 the storage icon to the right of the icon for the asset, for example, **DD**. The table in the right pane lists the backup copies.
5. Select one or more backup copies with the **Deletion Failed** or **Deletion Failed (Agent Catalog)** status from the table, and then click **Remove from PowerProtect**.

For backup copies with the **Deletion Failed (Agent Catalog)** status, click **Remove from PowerProtect** to remove the information from PowerProtect Data Manager for any backup copies that were successfully deleted from protection storage but for which the agent catalog was not deleted from the agent host.

The system displays a warning to confirm 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.

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 schedule is provided within the topic for each policy 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 currently assigned to an exclusion protection policy.
- **Disabled** – The asset is currently assigned to a disabled protection policy.
- **Protected** – The asset is currently assigned to an enabled protection policy.
- **Previously Protected** – The asset has been unassigned from a protection policy and has not yet been re-assigned 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](#) on page 95.

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 upon 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.

Export protection

This option enables you to export protection jobs and compliance records to a .CSV file so that you can download an Excel file of protection results data.

Steps

1. From the PowerProtect Data Manager UI, select **Protection > Protection Policies**.
The **Protection Policies** window appears, which displays the following information:
 - Asset type
 - Purpose
 - Group Name
 - Number of Protected Assets
 - Asset Capacity
 - Number of Failures
 - Number of SLA Violations
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 **Export**.
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.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 then click **Save**.

Disable a protection policy

From the PowerProtect Data Manager UI, you can disable a protection policy to temporarily stop running certain backup schedules 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 schedules of this policy, including synthetic full schedules, full schedules, and so on. Any replication, cloud tier, and extended retention schedules, 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. Protection jobs running for a disabled policy on page 100 provides information about jobs that continue to run when a policy is disabled.

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. Editing a protection policy on page 89 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 schedules 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 schedules 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 30. 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

Table 30. Protection jobs running when a policy is disabled (continued)

Job category	Purpose	Runs when policy is disabled?	System level overwrite?
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

1. From the PowerProtect Data Manager UI, select **Protection > Protection Policies**.
2. 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.
3. 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 schedules 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.

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.

Steps

- From the PowerProtect Data Manager UI, select **Protection > SLA Compliance**.
The **SLA Compliance** window appears.
- 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.
- 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.
 - Extended Retention**. If you choose this type, go to step 6.
 - Replication**. If you choose this type, go to step 7.
 - Cloud Tier**. If you choose this type, go to step 8.You can select only one type of Service Level Agreement.
- If you selected **Policy**, specify the following fields regarding the purpose of the new Policy SLA:
 - The **SLA Name**.
 - If applicable, select **Minimum Copies**, and specify the number of Backup, Replication, and Cloud Tier copies.
 - If applicable, select **Maximum Copies**, and specify the number of Backup, Replication, and Cloud Tier copies.
 - 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.
 - If applicable, select **Allowed in Cloud through Cloud Tier/Cloud DR**.
 - Click **Finish**, and then go to step 9.
- If you selected **Backup**, specify the following fields regarding the purpose of the new **Backup** SLA:
 - The **SLA Name**.
 - If applicable, select **Recovery Point Objective required (RPO)**, and then set the duration. The purpose of an RPO is business continuity planning, and 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.
 - 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.
 - e. If applicable, select **Retention Time Objective**, and specify the number of Days, Months, Weeks, or Years.
NOTE: For compliance validation to pass, the value set for the **Retention Time Objective** must match the lowest retention value set for the backup levels of this policy's target objectives. For example, if you set the synthetic full backup **Retain For** to 30 days but set the full backup **Retain For** to 60 days, the Retention Time Objective must be set to the lower value, in this case, 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 **Extended Retention**, specify the following fields regarding the purpose of the new Extended Retention 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, and indicates the maximum targeted period in which data (transactions) might be lost from an IT service due to a major incident.
NOTE: By default, the RPO provides a grace period of 1 day for SLA compliance verification. For example, with a weekly extended retention schedule, PowerProtect Data Manager provides 8 days for the RPO to pass the SLA Compliance verification.
 - 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.
 - d. If applicable, select **Retention Time Objective**, and specify the number of Days, Months, Weeks, or Years.
 - 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 **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.
 - d. If applicable, select **Retention Time Objective**, and specify the number of Days, Months, Weeks, or Years.
NOTE: For compliance validation to pass, the value set for the **Retention Time Objective** must match the lowest retention value set for the backup levels of this policy's target objectives.
 - 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.
 8. 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: For compliance validation to pass, the value set for the **Retention Time Objective** must match the lowest retention value set for the backup levels of this policy's target objectives.
 - d. If applicable, select the **Verify Retention Lock is enabled for all copies** option. This option is disabled by default.
 - e. Click **Finish**.
 9. 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**.
10. In the **Objectives** row of the **Summary** window, click **Edit**.
11. 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.

12. Click **Finish**.

An informational message appears to confirm that PowerProtect Data Manager has saved the protection policy.

13. 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 interval of time between the most recent backup of the asset and the compliance check is greater than the RPO requirement, then an alert indicates 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 interval of time between the most recent backup of the asset and the compliance check is less than the RPO requirement, the RPO of the asset is in compliance.

If multiple assets in a policy are out of compliance at the same time when a compliance check occurs, a single alert is generated and includes information for all assets that are out of compliance 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.

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

Export Asset Compliance

This option enables you to export compliance records to a CSV file 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 **Export Asset Compliance**.
The **Export Asset Compliance** 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**.

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.

When you define a protection rule, note the following requirements:

- 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.
- Virtual machine tags created in the **vSphere Client** can only be applied to a protection rule.
- To ensure the protection of homogeneous assets, the protection rule must specify a storage asset type.
- 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.

NOTE: Ensure that 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.

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.

Creating virtual machine tags in the vSphere Client

Creating virtual machine tags in the **vSphere Client** is supported by PowerProtect Data Manager with vSphere versions 6.5 and later. Tags enable you to attach metadata to the virtual assets in the vSphere inventory, which makes assets easier to sort and search for when creating a protection policy.

Asset inclusion in a PowerProtect Data Manager protection policy is based on the filtering criteria that you specify when creating a protection rule.

When you create a tag in the **vSphere Client**, the tag must be assigned to a category in order to group related tags together. When defining a category, you can specify the object types to which the tags will be applied and whether more than one tag in the category can be applied to an object. Within a single rule, you can apply up to 50 rule definitions to tags and categories, as shown in the following example where *Category* is the category name and *Bronze* is the tag name:

- Category:Category1,Tag:Bronze1
- Category:Category2,Tag:Bronze2
- Category:Category3,Tag:Bronze3
- ... Category:Category50,Tag:Bronze50

In the above example, category names and tag names that exceed 9 or 7 characters respectively reduce the limit for rule definitions in a single rule to less than 50. When rule definitions exceed the maximum limit, no virtual machines are backed up as part of the group, because no members are associated with the group. As a best practice, keep the number of rule definitions within a single rule to 10 or fewer and, in cases where there are a large number of rule definitions within a single rule, keep the number of characters in category or tag names to 10 or fewer.

To view existing tags for vCenter in the **vSphere Client**, select **Menu > Tags & Custom Attributes**, and then select the **Tags** tab. Click a tag link in the table to view the objects associated with this particular tag.

For PowerProtect Data Manager to include tagged assets in a protection rule based on the tags created for the vCenter, you must assign at least one tag to at least one virtual machine. Note that tags associated with containers of virtual machines (for example, a virtual machine folder) are not currently supported for tag associations to assets.

- NOTE:** Once virtual machines are associated with tags, the association is not reflected in the PowerProtect Data Manager UI until the timeout period has completed. The default timeout to fetch the latest inventory from the vCenter server is 15 minutes. When adding a protection rule and using tags as the asset filter, you must select **VM Tags**.

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. Click the tab to select the type of host for which you would like to add the protection rule, and then click **Add**. For example, **Virtual Machines**. The **Add Protection Rule** wizard opens to the **Select Protection Policy** page.
3. Select the target protection policy for the protection rule and then click **Next**. The **Asset Rules** 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**.
5. Define a protection rule:
 - a. Select an attribute. The available attributes depend on the selected host type and include names (such as **Datacenter Name** or **Host Name**), characteristics (such as **asset size**), tags (VM tags or namespace labels). The **Power State** attribute enables filtering of virtual machine hosts based on the state of the host (such as **Power On**, **Power Off**, or **Suspended**).
 - b. Select a matching criteria. The available matching criteria depend on the selected attribute:
 - For names, matching criteria include options such as **Begins with**, **Ends with**, **Contains**, **Does not contain**, **Equals**, **Match Regular Expression**, and **Does Not Match Regular Expression**.
The **VM Folder Name** and **VM Resource Pool** attributes support protection for all VM assets and resource pools in the selected folder and its subfolders.
 - For characteristics, matching criteria include options such as **Greater than or Less than**.
 - For tags, matching criteria include options such as **Includes**, **Does not include**, **In**, or **Not in**. The **In** and **Not in** criteria support multiple tags.
 - For **Power State**, matching criteria include options such as **Equals** and **Does Not Equal**.
 - Where the available matching criteria includes regular expressions, click  for a list of supported operators and effects in a separate dialog box.



NOTE:

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.

Regular expressions for all other attributes use Elasticsearch regex syntax. These expressions do not support predefined character classes.

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.

- 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.
The wizard adds a blank rule.
 - If you selected **AND**, the new rule appears in the same container.
 - If you selected **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.
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.
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, and file system 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 asset discovery

To schedule discovery in the PowerProtect Data Manager UI, complete the following steps:

Steps

1. Select **Infrastructure > Asset Sources**.
2. Select the **App/File System Host** tab.
3. Select the application host, 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.
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 on page 106 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.

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 rules changes still applies to the current release. For example, in PowerProtect Data Manager 19.4, 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 19.5 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.

Restoring Data and Assets

Topics:

- View backup copies available for restore
- Restoring a virtual machine or VMDK
- Restore an application-aware virtual machine backup
- 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 recovery within PowerProtect Data Manager display. Supported asset types include the following:

- **Virtual Machines**
- **File System**
- **Storage Group**
- **Kubernetes**

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 upon 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 virtual machine or VMDK

After virtual assets are backed up as part of a virtual machine protection policy in the PowerProtect Data Manager UI, you can perform image-level and file-level recoveries from individual or multiple virtual machine backups, and also restore individual virtual machine disks (VMDKs) to their original location.

PowerProtect Data Manager supports multiple data movers for restoring virtual machines, depending on the restore type and the vSphere capabilities. Restores are performed using one of the following data movers:

- **Transparent Snapshot Data Mover**—Starting in PowerProtect Data Manager version 19.9, Transparent Snapshot Data Mover (TSDM) is the default protection mechanism that is used for crash-consistent virtual machine policies when vCenter/ESXi version 7.0 U3 and later is installed in the environment. Review the section Prerequisites to restore a virtual machine on page 112 for specific restore type requirements for TSDM.
- **VADP**—VMware vStorage API for Data Protection (VADP) is the protection mechanism that is used for application aware virtual machine policies and crash-consistent policies that do not meet the TSDM software requirements. VADP is the only protection mechanism available in PowerProtect Data Manager versions 19.8 and earlier.
- **Storage vMotion** from protection storage to primary storage.

All types of recoveries are performed from the **Restore > Assets** window. Recovery options include the following:

- **Restore and Overwrite Original VM**: Restore to the original virtual machine.
- **Restore Individual Virtual Disks**: Restore select virtual disks to the original location.
- **Create and Restore to New VM**: Restore to a new virtual machine.
- **Instant Access VM**: Instant access to the virtual machine backup for browse and restore.
- **File Level Restore**: Restore individual files/folders the original or a new virtual machine
- **Direct Restore to ESXi**: Recover the virtual machine directly to an ESXi host without a vCenter server.

The **Restore** button, which launches the **Restore** wizard, is disabled until you select one or more virtual assets in the **Restore > Assets** window. Selecting multiple assets disables the **View Copies** button, since this functionality is available within the first page of the **Restore** wizard.

To access the **Restore and Overwrite Original VM**, **Create and Restore to New VM**, and **Instant Access VM** recovery types, or the **Restore Individual Virtual Disks** option, select one or more virtual assets and then click **Restore** to launch the **Restore** wizard.

To access the **File Level Restore** and **Direct Restore to ESXi** recovery options, select a virtual asset and then click **View Copies**.

In both instances, you must select a backup copy in the first page of the **Restore** wizard before you can go to the **Options** page, which displays the available recovery options.

NOTE: For all options, recovery in the PowerProtect Data Manager UI can only be performed if the backup or replica is on a DD system. If a replica backup does not exist on such storage, you must manually replicate this backup to DD storage before performing the restore.

The following sections describe each recovery option and provide instructions to perform the recovery.

NOTE: SQL virtual machine full database and transaction log restore from application-aware virtual machine protection policies must be performed using Microsoft application agent tools. The section Restore an application-aware virtual machine backup provides more information.

Restoring a virtual machine backup with the storage policy association

vSphere storage-based policies are used to communicate to the storage system details about how the virtual machine and its contents should be stored. At the time of backup, the existing policy assignments for the virtual machine will be stored in the backup copy.

During a restore to the original virtual machine in the PowerProtect Data Manager UI or the **vSphere Client**, you can select the **Restore Storage Policies** option if you want to restore any virtual machine disk-level or non-disk specific storage policy assignments.

This option is only applicable to virtual machine backup copies taken with PowerProtect Data Manager 19.6 and later. If you select this option but the virtual machine backup copy was created with PowerProtect Data Manager version 19.5 and earlier,

or the storage policy has been deleted from the vCenter Server, the virtual machine restore will proceed but any storage policy association will not be restored.

NOTE: Enabling this option requires vCenter version 6.7 and later.

Prerequisites to restore a virtual machine

Review the following requirements before you restore a virtual machine in PowerProtect Data Manager:

- Only the Administrator and the Restore Administrator roles can restore data.
- Ensure that you have added protection storage and the vCenter server, and that the protection of virtual machine copies has completed successfully.

To check, select **Infrastructure > Assets** and **Infrastructure > Asset Sources**.

- Ensure that protection of the virtual machines completed successfully. If the virtual machines have been backed up by a protection policy, the assets appear in the **Restore > Assets** window.
- If performing a restore to the original virtual machine, a minimum vCenter version 6.7 is required if you want to restore the virtual machine protection policy backup's storage policy assignments.
- If performing a restore to a new location, ensure that sufficient space is available on the target datastore.
- Verify that the virtual machine copy that is selected for restore has not expired.
- For restores of virtual machine protection policy backups using the Transparent Snapshot Data Mover (TSDM) protection mechanism, note the following:
 - For a **Restore to Original Folder and Overwrite Original Files**, the virtual machine must be currently protected by a policy that uses TSDM.
 - For a **Create and Restore to New VM**, the destination ESXi host where the new virtual machine will be created must have the vSphere Installation Bundle (VIB) installed and enabled.

Restore to the original virtual machine

A **Restore to Original Folder and Overwrite Original Files** recovers a virtual machine backup to its original location on the vCenter. This operation rolls the virtual machines that you backed up with the protection policy in PowerProtect Data Manager to an earlier point in time. Use this process for restoring the production system.

Prerequisites

Review Prerequisites to virtual machine restore before performing the restore.

About this task

NOTE: If the original virtual machine was deleted, a **Restore to Original Folder and Overwrite Original Files** recovery attempts to re-create the virtual machine. However, if the original virtual machine resources such as the datastore and cluster are no longer available, the restore fails and a **Restore to New** is required.

Steps

1. From the PowerProtect Data Manager UI, select **Restore > Assets**, and then select the **Virtual Machine** tab.
The **Restore** window displays all virtual machines available for restore.
2. Select the checkbox next to the appropriate virtual machines and click **Restore**.
You can also use the filter in the **Name** column to search for the asset name of the specific virtual machine or use the **File Search** button to search on specific criteria for files within backed-up virtual machines.
The **Restore** wizard appears.
3. On the **Select Copy** page, for each virtual machine that is listed in the table, select the radio button next to the virtual machine and click **Choose Copy**.
The **Choose Copy** dialog box appears.
 - NOTE:** If you click **Next** without choosing a copy, the most recent backup copy is used.
4. If the backup is on a DD system, click **DD**, and then select from one of the available copies that display in the table.
5. Click **OK** to save the selection and exit the dialog, and then click **Next**.

3. On the **Purpose** page, select **Restore Entire VMs** to restore the image-level virtual machine backup to the original location, and then click **Next**.

NOTE: If you specified any disk exclusions in the virtual machine protection policy, a message appears indicating that disks were excluded from this backup. If one of the excluded disks was a boot disk, the restore might not complete successfully.

The **Restore Type** page appears.

7. On the **Restore Type** page:

- a. Select **Restore to Original Folder and Overwrite Original Files**.

NOTE: If the system determines that the original virtual machine datastore(s) may be insufficient to complete the restore a warning is displayed. In this case, create more space in the original datastore(s), and then, select **Proceed Anyways**.

- b. Select the **Restore VM Tags** checkbox to restore vCenter tags and categories associated with this backup copy. Tags are backed up by default as part of the virtual machine protection policy backup.

NOTE: You can only select this option when restoring entire virtual machines. Any existing tags and categories on the assets in the restore location will be replaced with the tags and categories from the assets in the restored copy. If the tags and categories being restored do not exist in the vCenter Server at the time of the restore, or have been deleted, they will be re-created as part of the restore, along with the tag description and the cardinality settings that determine the relationship of tags within a category. If tags and categories on the vCenter have been renamed since the last backup, the renamed tags and categories will not be overwritten upon restore. For example, if a tag's ID is the same but the tag's name has been changed since the backup, a new tag is created based on the tag name in the backup copy being restored.

Upon successful restore, the replaced tags and categories will not be deleted in the **vSphere Client**, and can be viewed in the **Tags & Custom Attributes** window, or the **Tags** pane of the **Summary** window when the virtual machine is selected.

- c. Select **Restore Storage Policies** if you also want to restore any virtual machine disk-level or non-disk specific storage policy assignments.

If you select this option but the backup copy was taken with PowerProtect Data Manager 19.5 and earlier, or the storage policy is not available, the virtual machine restore will proceed but any storage policy association will not be restored.

NOTE: Enabling this option requires vCenter version 6.7 and later.

- d. For low-bandwidth environments, select **Enable DDBoost Compression**.

This option reduces network usage by compressing data on the protection storage system before transfer to the VM Direct Engine, which decompresses the data. Compression reduces restore times but increases CPU usage on both systems.

8. Click **Next**.

The **Networks** page appears if the virtual machine was backed up using PowerProtect Data Manager 19.9 or later. Otherwise, the **Options** page appears.

9. The **Networks** page displays the network adapters and associated networks the virtual machine had used when it was backed up. Click **Next** after reviewing this information and optionally performing one or both of the following actions.

NOTE: If a network used by an adapter is no longer accessible to the current virtual machine, a warning is displayed, and a different network should be selected for that adapter.

- a. To select a different network, click the associated drop-down control in the **Network** column, and then select an entry from the list.
- b. To change the initial power-on connection status of a network adapter, select or clear the associated check box in the **Connect at Power On** column.

If the current virtual-machine disk configuration is identical to the copy being restored, the **Summary** page appears, but if there is a mismatch, the **Options** page appears. This page displays the current configuration of the virtual machine along with any disks that have been added since the last backup.

10. On the **Options** page, for any hard disks in the current virtual machine configuration that were not part of the backup copy:

- Select **Delete disks that will be detached** to remove these disks upon restore.

- Clear **Delete disks that will be detached** to keep these disks in their original folders on the virtual machine after the restore. These disks will not be in the virtual machine configuration, but after the restore you can then use the **vSphere Client** to manually reattach or download these disks as appropriate.
11. Click **Next**.
The **Summary** page appears with a confirmation message indicating that the virtual machine will be powered off and that the virtual machine in the datastore will revert to the point in time of the selected backup copy before being powered back on.
 12. On the **Summary** page, click **Restore**.
An informational dialog box appears indicating that the restore has started.
 13. Go to the **Jobs** window to monitor the restore.
A restore job appears with a progress bar and start time.


Restore individual virtual disks

A **Restore Individual Virtual Disks** recovers individual virtual disks (VMDKs) to their original location on the vCenter, rolling the VMDKs that you backed up with the protection policy in PowerProtect Data Manager to an earlier point in time.

Prerequisites

Review Prerequisites to virtual machine restore before you perform the following procedure.

About this task

-  **NOTE:** When you restore individual VMDKs, only the selected disks are restored. The virtual machine configuration does not change.

Steps

1. From the PowerProtect Data Manager UI, select **Restore > Assets**, and then select the **Virtual Machine** tab.
The **Restore** window displays all virtual machines available for restore.
2. Select the checkbox next to the appropriate virtual machines and click **Restore**.
You can also use the filter in the **Name** column to search for the name of the specific virtual machine or click the **File Search** button to search on specific criteria.
The **Restore** wizard appears.
3. On the **Select Copy** page, for each virtual machine that is listed in the table, select the radio button next to the virtual machine and click **Choose Copy**.
The **Choose Copy** dialog box appears.
 **NOTE:** If you click **Next** without choosing a copy, the most recent backup copy is used.
4. If the backup is on a DD system, click **DD**, and then select from one of the available copies that display in the table.
5. Click **OK** to save the selection and exit the dialog, and then click **Next**.
6. On the **Purpose** page, select **Restore Individual Virtual Disks** to restore specific VMDKs.
7. For low-bandwidth environments, select **Enable DDBoost Compression**.
This option reduces network usage by compressing data on the protection storage system before transfer to the VM Direct Engine, which decompresses the data. Compression reduces restore times but increases CPU usage on both systems.
8. Click **Next**.
The **Select Disks** page displays.
9. From the **Backup Properties** pane, select the VMDKs that you want to restore, and then click **Next**. Note that individual VMDKs can only be restored to the original location.
The **Summary** page appears with a confirmation message indicating that the selected disk(s) will be overwritten in the current configuration with the copy from the backup.
10. On the **Summary** page, click **Restore**.
An informational dialog box appears indicating that the restore has started.
11. Go to the **Jobs** window to monitor the restore.
A restore job appears with a progress bar and start time.

Restore to a new virtual machine

A **Create and Restore to New VM** enables you to create a new virtual machine using a copy of the original virtual machine backup. Other than having a new name or location and a new vSphere VM Instance UUID, this copy is an exact replica of the virtual machine that you backed up with the protection policy in PowerProtect Data Manager.

Prerequisites

Review Prerequisites to virtual machine restore before you perform this procedure.

Steps

1. From the PowerProtect Data Manager UI, select **Restore > Assets**, and then select the **Virtual Machine** tab.
The **Restore** window displays all virtual machines available for restore.
2. Select the checkbox next to the appropriate virtual machines and click **Restore**.
You can also use the filter in the **Name** column to search for the name of the specific virtual machine or click the **File Search** button to run file level restore workflows on specific files within VMs.
The **Restore** wizard appears.
3. On the **Select Copy** page, for each virtual machine that is listed in the table, select the radio button next to the virtual machine and click **Choose Copy**.
The **Choose Copy** dialog box appears.
NOTE: If you click **Next** without choosing a copy, the most recent backup copy is used.
4. If the backup is on a DD system, click **DD**, and then select from one of the available copies that display in the table.
5. Click **OK** to save the selection and exit the dialog, and then click **Next**.
6. On the **Purpose** page:
 - Select **Restore Entire VMs** if you want to restore an image-level virtual machine backup.
NOTE: If you specified any disk exclusions in the virtual machine protection policy, a message appears indicating that disks were excluded from this backup. If one of the excluded disks was a boot disk, the restore might not complete successfully.
 - Select **Restore Individual Virtual Disks** if you want to restore only specific VMDKs.
NOTE: Individual disks can only be restored to the original location.
7. Click **Next**.
8. On the **Restore Type** page:
 - a. Select **Create and Restore to New VM**.
 - b. Select the **Restore VM Tags** checkbox to restore vCenter tags and categories associated with this backup copy. Tags are backed up by default as part of the virtual machine protection policy backup.
NOTE: You can only select this option when restoring entire virtual machines. Any existing tags and categories on the assets in the restore location will be replaced with the tags and categories from the assets in the restored copy. If the tags and categories being restored do not exist in the vCenter Server at the time of the restore, or have been deleted, they will be re-created as part of the restore, along with the tag description and the cardinality settings that determine the relationship of tags within a category. If tags and categories on the vCenter have been renamed since the last backup, the renamed tags and categories will not be overwritten upon restore. For example, if a tag's ID is the same but the tag's name has been changed since the backup, a new tag is created based on the tag name in the backup copy being restored.
Upon successful restore, the replaced tags and categories can be viewed in the **vSphere Client Tags & Custom Attributes** window, or the **Tags** pane of the **Summary** window when the virtual machine is selected.
 - c. For low-bandwidth environments, select **Enable DDBoost Compression**.
This option reduces network usage by compressing data on the protection storage system before transfer to the VM Direct Engine, which decompresses the data. Compression reduces restore times but increases CPU usage on both systems.
 - d. Click **Next**.
9. On the **VM Information** page:
 - a. From the **Restore to vCenter** list, select the vCenter server for the new virtual machine restore. This list displays any vCenter server that has been added from the **Assets** window.

When you select a vCenter server, available data centers appear.

- b. Select the destination data center.
 - c. Click **Next**.
10. On the **Restore Location** page:
 - a. Select the location within this data center that you want to restore the virtual machine by expanding the hierarchical view. For example, select a specific cluster, and then select a host within the cluster.
 - b. If you select an ESXi host within this page, the next page is unnecessary.
 - c. Click **Next**.
11. On the **ESX Host** page:
 - If you did not select a specific host in the previous step, select a host that is connected with the cluster, and then click **Next**.
 - If you selected a host in the previous step, this page indicates that a host is already selected and you can click **Next** to proceed.
12. On the **Datastore** page, select the datastore where you want to restore the virtual machine disks.
 - i** **NOTE:**


The **Total Estimated Space Needed for Recovery** is displayed and updated according to the specified disk provisioning type.

In the datastore list:

 - The free space in each datastore is displayed.
 - If a datastore is estimated to be smaller than required for recovery, it is displayed in red alongside an error icon.
 - Select **Browse...** to display the total capacity, provisioned capacity, and free capacity of all available datastore(s), and select a datastore.
 - a. If you are restoring multiple virtual machines, select the **Datastore** and **Provisioning Type** to use for all virtual machines.
 - b. If you are restoring one virtual machine:
 - To restore all disks to the same location, keep **Configure Per Disk** disabled, and select the datastore from the datastore list in the **Storage** column.
 - To restore disks to different locations, enable **Configure Per Disk**, and for each disk, select a datastore from the datastore list in the **Storage** column. Select how to provision the disk from the provisioning types in the **Disk Format** column.
 - i** **NOTE:** If you select a datastore whose estimated free space is smaller than required for recovery, a warning is displayed. In this case, you can select **Proceed Anyways** to continue, but it is recommended to create more space in the specified datastore(s) before doing so.
 - c. Click **Next**.
13. The **Networks** page appears if the virtual machine was backed up using PowerProtect Data Manager 19.9 or later. It displays the network adapters and associated networks the virtual machine had used when it was backed up. Click **Next** after reviewing this information and optionally performing one or both of the following actions.
 - i** **NOTE:** If a network used by an adapter is no longer accessible to the new virtual machine, a warning is displayed, and a different network should be selected for that adapter.
 - a. To select a different network, click the associated drop-down control in the **Network** column, and then select an entry from the list.
 - b. To change the initial power-on connection status of a network adapter, select or clear the associated check box in the **Connect at Power On** column.
14. On the **Options** page:
 - a. For **Select Access Level**, keep the slider set to **Yes** if you want to enable instant access for this restore.

When you select this option, the virtual machine is created and turned on while temporarily accessing the VMDKs from DD storage. Storage vMotion is initiated to the target datastore. The virtual machine becomes available for use when it is turned on.
 - b. [Optional] For the recovery options, select **Power on the virtual machine when the recovery completes** and **Reconnect the virtual machine's NIC when the recovery completes**. **Power on the virtual machine when the recovery completes** is selected by default when instant access is enabled.
 - c. Click **Next**.
15. On the **Summary** page, verify that the information you specified in the previous steps is correct, and then click **Restore**.





16. Go to the **Jobs** window to monitor the restore.

A restore job appears with a progress bar and start time. You can also click  next to the job to verify what steps have been performed, for example, when the instant access session has been created.


Instant access virtual machine restore

An **Instant Access VM** restore enables you to create a new virtual machine directly from the original virtual machine backup on protection storage for the purposes of instant backup validation and recovery of individual files. The instant access virtual machine is initially available for 7 days. This process does not copy or move any data from protection storage to the production datastore. An instant access virtual machine restore also provides the option to move the virtual machine to a production datastore when you want to retain access to the virtual machine for a longer time.

Steps

1. From the PowerProtect Data Manager UI, select **Restore > Assets**, and then select the **Virtual Machine** tab.
The **Restore** window displays all virtual machines available for restore.
2. Select the check box next to the appropriate virtual machines and click **Restore**.
You can also use the filter in the **Name** column to search for the name of the specific virtual machine, or click the **File Search** button to search on specific criteria.
The **Restore** wizard appears.
3. On the **Select Copy** page, for each virtual machine that is listed in the table, select the radio button next to the virtual machine and click **Choose Copy**.
The **Choose Copy** dialog box appears.
 **NOTE:** If you click **Next** without choosing a copy, the most recent backup copy is used.
4. If the backup is on a DD system, click **DD**, and then select from one of the available copies that display in the table.
5. Click **OK** to save the selection and exit the dialog, and then click **Next**.
6. On the **Purpose** page:
 - Select **Restore Entire VMs** if you want to restore an image-level virtual machine backup.
 **NOTE:** If you specified any disk exclusions in the virtual machine protection policy, a message appears indicating that disks were excluded from this backup. If one of the excluded disks was a boot disk, the restore might not complete successfully.
 - Select **Restore Individual Virtual Disks** if you want to restore only specific VMDKs.
 **NOTE:** Individual disks can only be restored to the original location.
7. On the **Restore Type** page:
 - a. Select **Instant Access VM**.
 - b. Select the **Restore VM Tags** checkbox to restore vCenter tags and categories associated with this backup copy.
 **NOTE:** You can only select this option when restoring entire virtual machines. Any existing tags and categories on the assets in the restore location will be replaced with the tags and categories from the restored copy. If the tags and categories being restored do not exist in vCenter at the time of the restore, or have been deleted, they will be re-created as part of the restore, along with the tag description and the cardinality settings that determine the relationship of tags within a category. If tags and categories on the vCenter have been renamed since the last backup, the renamed tags and categories will not be overwritten upon restore. For example, if a tag's ID is the same but the tag's name has been changed since the backup, a new tag is created based on the tag name in the backup copy being restored.
Upon successful restore, the replaced tags and categories can be viewed in the **vSphere Client Tags & Custom Attributes** window, or the **Tags** pane of the **Summary** window when the virtual machine is selected.
 - c. Click **Next**.
8. On the **VM Information** page:
 - a. Select whether you want to use the original virtual machine name for the instant access virtual machine restore, or rename the instant access virtual machine by appending a suffix to the original name.
 - b. From the **Restore to vCenter** list, select the vCenter server for the instant access virtual machine restore. You can select the vCenter of the original virtual machine backup, or another vCenter. This list displays any vCenter server that has been added from the **Assets** window.

When you select a vCenter server, available data centers appear.

- c. Select the destination data center.
 - d. Click **Next**.
9. On the **Restore Location** page, select the location within this data center that you want to restore the virtual machine by expanding the hierarchical view. For example, select a specific cluster, and then select a host within the cluster. If you select an ESXi host within this page, the next page is unnecessary. Click **Next**.
10. On the **ESX Host** page:
- If you did not select a specific host in the previous step, select a host that is connected with the cluster, and then click **Next**.
 - If you selected a host in the previous step, this page indicates that a host is already selected and you can click **Next** to proceed.
11. The **Networks** page appears if the virtual machine was backed up using PowerProtect Data Manager 19.9 or later. It displays the network adaptors and associated networks the virtual machine had used when it was backed up. Click **Next** after reviewing this information and optionally performing one or both of the following actions.
- NOTE:** If a network used by an adaptor is no longer accessible to the new virtual machine, a warning is displayed, and a different network should be selected for that adaptor.
- a. To select a different network, click the associated drop-down control in the **Network** column, and then select an entry from the list.
 - b. To change the initial power-on connection status of a network adaptor, select or clear the associated check box in the **Connect at Power On** column.
12. On the **Options** page:
- a. Specify a name for the instant Access virtual machine.
 - b. Optionally, select **Power on the virtual machine when the recovery completes** and **Reconnect the virtual machine's NIC when the recovery completes**. **Power on the virtual machine when the recovery completes** is selected by default for instant access virtual machine restores.
 - c. Click **Next**.
13. On the **Summary** page, verify that the information you specified in the previous steps is correct, and then click **Restore**. A confirmation message displays indicating that the restore has been initiated and providing the option to go to the **Jobs** window to monitor the restore progress.
14. Go to the **Jobs** window to view the entry for the instant access virtual machine recovery and verify when the recovery completes successfully. You can also click  next to the job to verify what steps have been performed, for example, when the instant access session has been created.

Results

To monitor and manage the instant access virtual machine recovery, select **Restore > Running Sessions**, and then click the **Instant Access** tab. From this window, you can also extend the instant access virtual machine session beyond the default period of 7 days.

- NOTE:** On a single-node protection storage system such as a DD system, instant access/restore functionality has been enhanced to return a failure message when overwhelmed with traffic. For example, if on the target node or the ESXi host there are Live VM and/or Instant Restore sessions that are in conflict, instant access/restore jobs will fail with a message indicating a resource contention issue. If this occurs, you need to clear the conflicts and then restart the session in order for the job to execute.

Manage and monitor Instant Access sessions

In the PowerProtect Data Manager UI, the **Instant Access** tab of the **Restore > Running Sessions** window enables you to monitor vMotion events, and to manage the status of a virtual machine restore to new or instant access-virtual machine restore. For example, you can extend the availability period or delete an instant access virtual machine.

- NOTE:** The Instant Access Sessions that are used by a SQL application-aware self-service restore are displayed in the PowerProtect Data Manager UI, but management is disabled. Use the SQL application-aware self-service restore UI to manage these sessions.

When the Jobs window indicates that a recovery has completed successfully, go to **Restore > Running Sessions > Instant Access** to access information about the sessions. This window enables you to monitor and manage all exported copies that you have created from protection storage. An active restore session with a state of **Mounting** indicates that the restore is still in

progress. Once the state changes to **Mounted**, the restore is complete and the instant access virtual machine is ready. When you select the session in the table, you can choose from three options:

- **Extend** —Click to extend the number of days the instant access virtual machine restore is available. The default retention period of an instant access virtual machine restore is 7 days.
- **Migrate** —Click to open the **Migrate Storage vMotion** wizard, which enables you to move the instant access virtual machine to a protection datastore. Migrate an instant access session provides instructions.
- **Delete** —Click if you no longer require the active restore session. Note that you can also vMotion from inside the vCenter server, and PowerProtect Data Manager removes the Instant Access Session upon detection.

For instant access virtual machine restores, availability of the instant access virtual machine session is also indicated in the **vSphere Client**. The session appears in the **Recent Tasks** pane, and you can expand the cluster and select the instant access virtual machine to view summary information, as shown in the following figure.




Figure 3. instant access virtual machine restore in the vSphere Client

Migrate an Instant Access session

Once you validate that the instant access virtual machine is the virtual machine that you require for production, click **Migrate** to open the **Migrate Storage vMotion** wizard, which enables you to select the session and move the virtual machine to a production datastore.

Steps

1. From the PowerProtect Data Manager UI, select **Restore > Running Sessions**, and then click the **Instant Access** tab.
2. Select a session from the table that is in **Mounted** state, and click **Migrate**. The **Migrate Storage vMotion** wizard displays.
3. On the **Disk Files Datastore** page, select the datastore where you want to relocate the instant access virtual machine, and then click **Next**.
 - To migrate all VMDKs to the same datastore, keep the **Configure per disk** slider to the left, and then select the datastore from the **Storage** list.
 - To migrate VMDKs to separate datastores, move the **Configure per disk** slider to the right, and then:
 - a. Select a datastore for each disk from the **Storage** list.
 - b. Select the type of provisioning you want to apply to the disk from the **Disk Format** list.
4. On the **Summary** page, review the information to ensure that the details are correct, and then click **Migrate**.
5. Go to the **Jobs** window or the **Instant Access** window to view the progress of the migration.

In the **Jobs** window, the migration job appears with a progress bar and start time. You can also click  next to the job to verify what steps have been performed. In the **Instant Access** window, you can monitor the vMotion status of the

migration. When a vMotion is in progress, the status indicates **VMotioning**. Once the storage vMotion for the session is complete, the status of the session changes to **Deleting** as the session is being removed from the **Instant Access** window.

File level restore to original virtual machine

A file level restore to original virtual machine enables you to recover individual files from backups of virtual machines or VMDKs performed in PowerProtect Data Manager to the same or a new location on the original vCenter Server. Only the Administrator and the Restore Administrator roles can restore data.

Prerequisites

- Review the section Supported platform versions for file-level restore for supported platform and operating system versions.
- Review the section File-level restore and SQL restore limitations on page 247.

NOTE: For file-level restores, you can only restore files:

- From a Windows backup to a Windows machine, or from a Linux backup to a Linux machine.
- To virtual machines within the same vCenter.

Steps

1. From the PowerProtect Data Manager UI, select **Restore > Assets**, and then select the **Virtual Machine** tab.
The **Restore** window displays all the virtual machines available for restore.
2. Select the checkbox next to the virtual machine that you want to recover from, and then click **View Copies**.
You can also use the filter in the **Name** column to search for a specific virtual machine name, or click the **File Search** button to search on specific criteria.

The **Restore > Assets** window provides a map view in the left pane and copy details in the right pane.

When a virtual machine is selected in the map view, the virtual machine name displays in the right pane with the copy locations underneath. When you select a location in the left pane, for example, a DD system, the copies on that system display in the right pane.

3. If the backup is on a DD system, click **DD**, and then select from one of the available copies that display in the table.
4. In the right pane, select the checkbox next to the virtual machine backup you want to restore, and then click **File Level Restore**.

The **File Level Recover** wizard appears.

5. On the **Restore Type** page, select **Restore to Original Virtual Machine**, and then click **Next**.
6. On the **Mount Copy** page:

- a. To initiate the disk mount, type the guest operating system user credentials:
 - If there are administrator-level credentials associated with the virtual assets or protection policy being restored, specify end-user credentials.
 - If there are no administrator-level credentials associated with the virtual assets or protection policy being restored, specify administrator credentials. These credentials will be handled as end-user credentials.
- b. (Optional) Leave **Keep FLR Agent Installed** selected when you want the **FLR Agent** to remain on the destination virtual machine after the restore completes.
- c. Click **Start Mount** to initiate the disk mount. A progress bar indicates when the mount completes.

NOTE: You cannot browse the contents of the virtual machine backup until the mounting of the destination virtual machine completes successfully.

When validated, the **FLR Agent** is installed automatically on the restore destination, if it is not already installed. The **FLR Agent** facilitates the mounting and unmounting of disks and the browsing of files in the destination virtual machine and the backup copy. In order to complete the automatic **FLR Agent** installation, on Windows virtual machines the user must be an administrator account, and on Linux virtual machines the user must be the root user account, or a user in the operating system's local sudoers list. The section **FLR Agent for virtual machine file-level restore** on page 248 provides more information.

- d. Upon successful mount, click **Next**.
7. On the **Select Files to Recover** page:
 - a. Expand individual folders to browse the original virtual machine backup, and select the objects that you want to restore to the destination virtual machine.
 - b. Click **Next**.

NOTE: When you browse for objects to recover on this page, each directory or hard drive appears twice. As a result, when you select an object from one location, the object is selected in the duplicate location as well.

8. On the **Options** page, select from one of the following options, and then click **Next**.
 - Restore to Original Folder and Overwrite Original Files—Select this option to restore all selected files to their original location on the original virtual machine.
 - Restore to an Alternate Folder—Select this option if you want to restore to a new folder in a new location on the original virtual machine.
9. On the **Summary** page:
 - a. Review the information to ensure that the restore details are correct. You can click **Edit** next to any row to change the information.
 - b. Click **Restore**.
10. Go to the **Jobs** window to monitor the restore.

A restore job appears with a progress bar and start time.

File level restore to alternate virtual machine

A file level restore to alternate virtual machine enables you to recover individual files from backups of virtual machines or VMDKs performed in PowerProtect Data Manager to a new location on a new virtual machine. This restore can be performed to a primary or secondary vCenter Server. Only the Administrator and the Restore Administrator roles can restore data.

Prerequisites

- Review the section Supported platform versions for file-level restore for supported platform and operating system versions.
- Review the section File-level restore and SQL restore limitations on page 247.

NOTE: For file-level restores, you can only restore files:

- From a Windows backup to a Windows machine, or from a Linux backup to a Linux machine.
- To virtual machines within the same vCenter.

Steps

1. From the PowerProtect Data Manager UI, select **Restore > Assets**, and then select the **Virtual Machine** tab.

The **Restore** window displays all the virtual machines available for restore.
2. Select the checkbox next to the virtual machine that you want to recover from, and then click **View Copies**.

You can also use the filter in the **Name** column to search for a specific virtual machine name, or click the **File Search** button to search on specific criteria.

The **Restore > Assets** window provides a map view in the left pane and copy details in the right pane.

When a virtual machine is selected in the map view, the virtual machine name displays in the right pane with the copy locations underneath. When you select a location in the left pane, for example, a DD system, the copies on that system display in the right pane.
3. If the backup is on a DD system, click **DD**, and then select from one of the available copies that display in the table.
4. In the right pane, select the checkbox next to the virtual machine backup you want to restore, and then click **File Level Restore**.

The **File Level Recover** wizard appears.
5. On the **Restore Type** page, select **Restore to Alternate Virtual Machine**, and then click **Next**.
6. On the **Select Target VM** page, choose from one of the following options:
 - Search for a target virtual machine by typing the name.
 - Browse from the available vCenter Servers to locate the destination virtual machine.
7. On the **Mount Copy** page:
 - a. To initiate the disk mount, type the guest operating system user credentials:
 - If there are administrator-level credentials associated with the virtual assets or protection policy being restored, specify end-user credentials.
 - If there are no administrator-level credentials associated with the virtual assets or protection policy being restored, specify administrator credentials. These credentials will be handled as end-user credentials.
 - b. (Optional) Leave **Keep FLR Agent Installed** selected when you want the **FLR Agent** to remain on the destination virtual machine after the restore completes.

- c. Click **Start Mount** to initiate the disk mount. A progress bar indicates when the mount completes.

i **NOTE:** You cannot browse the contents of the virtual machine backup until the mounting of the destination virtual machine completes successfully.

When validated, the **FLR Agent** is installed automatically on the restore destination, if it is not already installed. The **FLR Agent** facilitates the mounting and unmounting of disks and the browsing of files in the destination virtual machine and the backup copy. In order to complete the automatic **FLR Agent** installation, on Windows virtual machines the user must be an administrator account, and on Linux virtual machines the user must be the root user account, or a user in the operating system's local sudoers list. The section **FLR Agent for virtual machine file level restore** on page 248 provides more information.

- d. Upon successful mount, click **Next**.
8. On the **Select Files to Recover** page:
 - a. Expand individual folders to browse the original virtual machine backup, and select the objects that you want to restore to the destination virtual machine.
 - b. Click **Next**.

i **NOTE:** When you browse for objects to recover on this page, each directory or hard drive appears twice. As a result, when you select an object from one location, the object is selected in the duplicate location as well.
9. On the **Restore Location** page:
 - a. Browse the folder structure of the destination virtual machine to select the folder where you want to restore the objects.
 - b. Click **Next**.
10. On the **Summary** page:
 - a. Review the information to ensure that the restore details are correct. You can click **Edit** next to any row to change the information. If you are not restoring to the original virtual machine, an additional field appears for the **Target VM**.
 - b. Click **Restore**.
11. Go to the **Jobs** window to monitor the restore.
A restore job appears with a progress bar and start time.

Direct restore to ESXi

If the virtual machine you protected with PowerProtect Data Manager was a vCenter virtual machine, but this virtual machine and vCenter is now lost or no longer available, direct restore to ESXi enables you to recover the virtual machine directly to an ESXi host without a vCenter server.

Prerequisites

Direct Restore to ESXi restore requires either the embedded VM Direct engine with PowerProtect Data Manager, or an external VM Direct appliance that is added and registered to PowerProtect Data Manager.

Additionally, ensure that you disconnect the ESXi host from the vCenter server.

Steps

1. From the PowerProtect Data Manager UI, select **Restore > Assets**, and then select the **Virtual Machine** tab.
The **Restore** window displays all of the virtual machines available for restore.
2. Select the checkbox next to the desired virtual machine and click **View Copies**.

i **NOTE:** If you cannot locate the virtual machine, you can also use the filter in the **Name** column to search for the name of the specific virtual machine or click the **File Search** button to search on specific criteria.

The **Restore > Assets** window provides a map view in the left pane and copy details in the right pane.

When a virtual machine is selected in the map view, the virtual machine name displays 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 display in the right pane.

3. If the backup is on a DD system, click **DD**, and then select from one of the available copies that display in the table.
4. In the right pane, select the checkbox next to the virtual machine backup you want to restore, and then click **Direct Restore to ESXi**.
The **Direct Restore to ESXi** wizard appears.
5. On the **Options** page:

- a. (Optional) Select **Reconnect the virtual machine's NIC when the recovery completes**, if desired. **Power on the virtual machine when the recovery completes** is selected by default.
 - b. For low-bandwidth environments, select **Enable DDBoost Compression**.
This option reduces network usage by compressing data on the protection storage system before transfer to the VM Direct Engine, which decompresses the data. Compression reduces restore times but increases CPU usage on both systems.
 - c. Click **Next**.
6. On the **ESX Host Credentials** page:
 - a. In the **ESX Host** field, type the IP of the ESXi server where you want to restore the virtual machine backup.
 - b. Specify the root **Username** and **Password** for the ESXi Server.
 - c. Click **Next**.
 7. On the **Datastore** page, select the datastore where you want to restore the virtual machine disks, and then click **Next**.
 - To restore all of the disks to the same location, keep the **Configure per disk** slider to the left, and then select the datastore from the **Storage** list.
 - To restore disks to different locations, move the **Configure per disk** slider to the right, and then:
 - a. For each available disk that you want to recover, select a datastore from the **Storage** list.
 - b. Select the type of provisioning you want to apply to the disk from the **Disk Format** list.
 8. On the **Summary** page:
 - a. Review the information to ensure that the details are correct.
 - b. Click **Restore**.
 9. Go to the **Jobs** window to monitor the restore.
A restore job appears with a progress bar and start time.

Restore an application-aware virtual machine backup

When virtual machine applications are protected within a protection policy in PowerProtect Data Manager, you can recover the application data using the Microsoft application agent, or perform a centralized restore within the PowerProtect Data Manager UI.

The *PowerProtect Microsoft Application Agent SQL Server User Guide* provides instructions on how to restore an application-aware virtual machine using the VM Direct SQL Server Management Studio (SSMS) plug-in.

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.

- Specify the following storage information:
 - DD system IP where the recovery backups are stored.
 - DD NSF Export Path where the recovery backups are stored.
 - Click **Connect**.
- Select the PowerProtect Data Manager instance that you would like to restore, and then click **OK**.
- Select the backup file that you would like to use for recovery, and then click **Recover**.
- 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

- From the PowerProtect Data Manager UI, select **Infrastructure > Assets**.
- On the **Assets** window, select the tab that contains the asset you want to recall from Cloud Tier, and then click **View Copies**.
- Click **DD**, and then select from one of the available copies that appear in the table.
- Click **Recall**.
The **Recall from Cloud** dialog box appears.
- In the **Retain until** box, specify how long you want to keep the copy on the active tier, and then click **OK**.
- 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**.

i **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:

- Managing system backups for server disaster recovery
- Prepare the DD system recovery target (NFS)
- Configure PowerProtect Data Manager server DR backups
- Record settings for server DR
- Manage PowerProtect Data Manager server DR backups
- Restore PowerProtect Data Manager from server DR backups
- Troubleshooting NFS backup configuration issues
- Troubleshoot recovery of PowerProtect Data Manager
- Quick recovery
- Recover a failed PowerProtect Data Manager backup

Managing system backups for 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 system disaster recovery (DR) backups.

Each backup is considered a full backup although it is created in an incremental manner. The persistent data that is saved in a backup includes the lockbox and ElasticSearch databases. The backup operation creates a point-in-time snapshot of the database while the system is in a quiesced state. While the system is quiesced, user functionality is limited. After the snapshot completes and while PowerProtect Data Manager copies the snapshots to protection storage, full user functionality is restored.

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 on page 129 provides instructions.

File Search indexes are backed up for DR recovery along with other component DR backups.

You can back up to only one protection storage system at a time. When you specify a new protection storage system for backup, you overwrite the existing protection storage system selection. If you have more than one protection storage system, you can change which protection storage system holds the server DR backup.

Server DR protection storage types

PowerProtect Data Manager supports two types of protection storage for server DR: NFS and DD Boost.

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.

NFS

NFS is the legacy storage type for PowerProtect Data Manager server DR. 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.

Starting with PowerProtect Data Manager 19.9, NFS storage is deprecated.

DD Boost

DD Boost is the recommended storage type for PowerProtect Data Manager server DR. 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 password. Changes to the predefined administrator account password prompt corresponding updates to the DD Boost user password. Recovery from server DR backups requires the predefined administrator account 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 on page 38 provides instructions.

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.

The *PowerProtect Data Manager Cloud Disaster Recovery Administration and User Guide* provides more information about Cloud DR workflows within PowerProtect Data Manager.

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.

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**.
3. 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.
 - ① **NOTE:** For an external DD system, specify a path similar to the following, `/data/coll1/<path>`, where `<path>` is the MTree that stores the system backups.
4. When the progress message indicates that the save operation is complete, click **Close**.
5. In the **Summary** tab in the **Protocols** pane, click **NFS Exports**.
6. Under **NFS Protocols > Exports**, select the DD MTree from the list of exports and click **Add Clients**.
7. In the **Add Clients** window, provide the following information, and then click **OK**.
 - **Client**—IP address or hostname of the PowerProtect Data Manager.

NOTE: To configure DR protection for an existing Search cluster, add the IP address or hostname of the Search cluster to the NFS client list.

- Accept the default settings for the rest of the fields.
- **Current Selection**—Ensure that the list includes `no_root_squash`, which is required for permission for PowerProtect Data Manager to change the directory structure on the NFS share.

Configure PowerProtect Data Manager server DR backups


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)* on page 127.

If you plan to use DD Boost for protection storage, add the DD system as protection storage. *Protection storage* on page 38 provides instructions.

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. Select **Enable backup**.
4. For NFS, configure the backup with the following attributes:
 - a. For **Protocol**, select **NFS**.
 - b. In the **PowerProtect DD System** field, type the IP address or hostname of the DD system for the backup.
 - c. In the **NFS Export Path** field, type the NFS path where server DR backups are stored on the target DD system.
5. For DD Boost, configure the backup with the following attributes:
 - a. For **Protocol**, select **DDBoost**.
 - b. From the **PowerProtect DD System** drop-down list, select an existing protection storage system. 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.
6. Configure the backup frequency and duration:
 - a. Type an interval between server DR backups, in hours. This setting controls backup frequency, and the allowed values are 1 to 24 hours.
 - b. Type the number of days for which PowerProtect Data Manager should retain server DR backups. The allowed values are 2 to 30 days.
7. 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.


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 certain information to recover your systems. Record the following information on a local drive outside PowerProtect Data Manager:

Steps

1. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.
2. Record the PowerProtect Data Manager build number.
Customer Support can provide this information, which is not mandatory.
3. Record the PowerProtect Data Manager hostname.
4. 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.
5. Click , select **Disaster Recovery**, and then click **Configuration**.
6. Record whether server DR storage uses NFS or DD Boost.
7. If you use NFS for server DR storage, record the protection storage system IP address or hostname, and the NFS export path.
8. If you use DD Boost for server DR storage, record which protection storage system stores the server DR backups.
Record the protection storage system IP address and hostname or FQDN.
9. Use the REST API to record additional configuration information:

Run the GET /Configurations API (api/v2/configurations) from PowerProtect Data Manager and save the details for network information.

- a. Obtain a PowerProtect Data Manager authentication token:

```
curl --request POST
'https://<ppdm_ip>:8443/api/v2/login' --header
'Content-Type: application/json' --data
'{"username":<user>,"password":<password>}' -k
```

- b. Use the authentication token to get the configuration from PowerProtect Data Manager:

```
curl --request GET
'https://<ppdm_ip>:8443/api/v2/configurations' --header
'Content-Type: application/json' --header
'Authorization: Bearer <token>' -k
```

Manage PowerProtect Data Manager server DR backups

View PowerProtect Data Manager server DR backups and perform manual backups. You can view the last 5 server DR 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 cannot manually delete scheduled backups. However, you can delete manual backups.

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 `UserDR-`. 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 the arrow icon.
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**.

Restore PowerProtect Data Manager from server DR backups

You can restore PowerProtect Data Manager from a server DR backup on a protection storage system.

Prerequisites

- Ensure that all the information listed in Record settings for server DR on page 129 is available.
- Ensure that the FQDN of the PowerProtect Data Manager is the same as the host name.
- Ensure that the VM for PowerProtect Data Manager is powered on.
- Ensure that you have set up the recovery target system. See Prepare the DD system recovery target (NFS) on page 127.

To restore 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.

- NOTE:** If the Search Engine nodes from the previous PowerProtect Data Manager installation are still hosted on the vCenter, delete the Search Engine nodes from the vCenter before you restore the PowerProtect Data Manager system. The disaster recovery process redeploys the Search Engine nodes as part of the restore operation.

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 is on a different FQDN, see Troubleshoot recovery of PowerProtect Data Manager on page 133.

If a Search Engine is present in the recovery backup when you restore the PowerProtect Data Manager system, the Search Engine is automatically restored.

Steps

1. Use the OVA file to deploy a new PowerProtect Data Manager system.
2. On the **Install** window under **Welcome**, select **Restore Backup**.
3. (Optional) To keep the PowerProtect Data Manager server in recovery mode after the restore completes, select the checkbox.

When this option is enabled, PowerProtect Data Manager enters into recovery mode and stops scheduled workflows from running.

4. To restore 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.
5. To restore 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 PowerProtect Data Manager system.
 - d. Type the password for the predefined administrator account.
 - e. Click **Connect**.
A list of the available recovery backups appears.
6. Select the backup from which to recover the system, and then click **Recover**.
The recovery starts. Recovery can take a few minutes.

Results

When recovery is complete, the PowerProtect Data Manager login page appears.

When you log in to PowerProtect Data Manager, if the option to keep the PowerProtect Data Manager server in recovery mode was selected, 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. If you want to return PowerProtect Data Manager to full operational mode and enable scheduled workflows, click **Return to full operational mode**.

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.

Recovering 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**.

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. Locate the backup manifest file:
 - a. Connect to the PowerProtect Data Manager console as an admin user.
 - b. Browse the directory path `/data01/server_backups/<PPDM Hostname>_<NodeID>`.
 - c. Run `grep -Rnwa -e '<Name>' --include=*.manifest`
3. Open the backup manifest file.
4. Locate the Components section, which contains Search Cluster.
The values for the following fields that are listed in the Search Cluster section are needed for the POST call in the next step.
 - `Name=Id`
 - `BackupPath`, which contains `<NFSEHost>:/data/coll/<NFSEExportFolder>/<NFSEDirPath>/SearchCluster`

For example:

```
"Components": [
  { "name": "SearchCluster",
```

```
{
  "id": "c25290d9-a88c-4a15-9e7c-656f186209ae",
  "version": "v2",
  "backupPath": "10.25.12.74:/data/coll/serverdr_backup/vm-qa-0091_6ce36793-3379-45d2-84bd-d8bde69e52d4/SearchCluster",
  "backupStatus": "SUCCESSFUL",
  "backupsEnabled": true
}
```

where:

- o NFSHost = "10.25.12.74"
- o NFSExport = "/data/coll/serverdr_backup"
- o NFSDirPath = "vm-qa-0091_6ce36793-3379-45d2-84bd-d8bde69e52d4/SearchCluster"
- o Name = "c25290d9-a88c-4a15-9e7c-656f186209ae"

5. Use the REST API to run the following POST call:

```
https://<PPDM_IP>:8443/api/v2/search-clusters/component-backups/
<Name>/restore

{
  "ddDirectoryPath" : "<NFSDirPath>",
  "ddHost" : "<NFSHost>",
  "ddNfsExportName" : "<NFSExport>"
}
```

6. 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**.

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 DD Boost storage unit.

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://PPDM_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.

Quick recovery

After a disaster, the quick recovery feature allows you to restore assets and data that you replicated to a destination system at a remote site.

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* on page 134 shows the initial configuration with both sites replicating copies to each other. Figure *Separate datacenters, after disaster* on page 135 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 4. Separate datacenters, before disaster




Figure 5. 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 6. 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:



- Attach at least two protection storage system 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.

Before you use the quick recovery remote view, add the destination system to the list of remote systems on the source.

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.
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 from the list.
Alternatively, you can click **Add Credentials** from this list to add new credentials. 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 on page 60 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. Dell Technologies recommends that you perform a test restore.

Metadata syncs between source and destination systems every six hours. If backups are not visible, allow sufficient time for the first sync before troubleshooting.

Edit a remote system

You can change the descriptive name of the remote system, as well as the REST API port number and credentials. 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**.
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 contains an additional **Remote Systems**  icon in the banner.

When you click **Remote Systems**, PowerProtect Data Manager presents a drop-down that contains the names 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.

Recover a failed PowerProtect Data Manager backup

Steps

1. Redeploy the PowerProtect Data Manager OVA.
2. Contact Customer Support.

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

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 on page 163 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 the notification category.
6. From the **Severity** list, select the notification severity.
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.
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**.
The **Alerts** window displays alert information in a table. You can filter the alerts by Severity, Date, Category, or Acknowledge.
2. Select the **System** tab.
The **System** tab displays all alert types.
3. To view more details about a specific entry, click  next to the entry in the table.
4. 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.
5. To acknowledge the alert, select the alerts and then click **Acknowledge**.
6. To add or edit a note for the alert, click **Add/Edit Note**, and when finished, click **Save**.
7. To export a report of alert information to a .CSV file which you can download for Excel, select an entry in the table and then click **Export**.

 **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**.
 **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

Use the **Protection Jobs** and **System Jobs** windows in the PowerProtect Data Manager UI to monitor the status of certain data protection, system, and maintenance jobs and to 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. Jobs are categorized as protection jobs or system jobs.

You can also view details for a job group and individual jobs and tasks. 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, so that you can monitor the status of individual jobs and tasks, view job and task details, and perform certain operations on jobs and tasks.

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 on page 146 provides more information.

- ① **NOTE:** The **Protection Jobs** and **System Jobs** windows have been optimized for a screen resolution of at least 1920 x 1080 pixels with 100% scaling. Display issues might occur for smaller screens. Set your screen resolution to at least 1920 x 1080 pixels with 100% scaling.

Monitor and view jobs

Use the **Protection jobs** and **System jobs** windows to monitor and view status information for PowerProtect Data Manager operations.

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

① **NOTE:** This job type does not apply to SAP HANA databases.

- Export Reuse
- Protect
- Replicate
- Restore

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.

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 31. 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
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.
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 on page 146](#) 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 7. List of protection jobs

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 will be shorter than the duration indicated in the Job Metrics.

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 SQL 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 the following:

- Application agent assets from a protection job that was performed for an application agent previous to version 19.7. To display the correct information for these assets, update the application agent on these assets to the current version.
- 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.
- **Stats 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 will appear in the **Job ID Summary** window of every asset type.

Table 32. 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:

Table 32. Job ID Summary window details (continued)

Column	Description
	<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.

- 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 for the job or task, and indicates the amount of time that was required to complete each step.
- **Details**—Displays statistics and summary information, such as the start time and end time, asset size, duration, and so forth.
- **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 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

- From the PowerProtect Data Manager UI left navigation pane, select **Jobs > System Jobs**.
- 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 on page 146 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.

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 for the job or task and indicates the amount of time that was required to complete each step.
- **Details**—Displays statistics and summary information, such as the start time and end time, asset size, duration, and so forth.
- **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** and **System Jobs** windows provide options to filter, group, and sort the information that appears. Select a job to display its **Job ID Summary** window.

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**
- **Completed**

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. 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

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 and Exchange 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 database	SQL Host
	SQL Instance
Oracle database	Oracle Host

Table 34. Group by options (continued)


Asset type	Options
	Oracle Instance
File System	File System Host
	File System Host OS
Microsoft Exchange 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 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.	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 36. Job ID Summary window for protection jobs

Filtering options	Sorting options
	① 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.

① **NOTE:**

- If a policy with both protection and Cloud Data Recovery (CDR) stages fails, the CDR job is canceled and cannot be restarted.
- Cloud Native Entity jobs cannot be restarted.

Steps

1. From the PowerProtect Data Manager UI left navigation pane, select **Jobs > Protection Jobs** or **Jobs > System Jobs**. The **Protection Jobs** or **Systems Jobs** window appears, displaying 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**.
3. To restart a failed 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**.

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, `-jar /$(APP_NAME)/lib/workflow-manager.jar`, 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

Upon configuration, the workflow service is scheduled to run every 30 minutes to determine if any jobs or tasks have failed. If a restart occurred, the status will indicate **Running** or **Queued**. To view whether a failed job or task has been 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 PostgreSQL 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:

① **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 \
```

① **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`

① **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/ppdac-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:

① **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

① **NOTE:** The **Cancel** operation is available for the following supported jobs and tasks only:

- Backup and restore of:
 - Virtual machine

- Kubernetes
- NAS
- File System agent
- SQL agent
- Server DR
- Cloud DR
- Backup (only) of:
 - Exchange agent
 - Oracle agent
 - SAP HANA agent
- 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** or **Jobs > System Jobs**. The **Protection Jobs** or **Systems 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**.

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 **Protection Jobs** or **System Jobs** 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. 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.

Steps

1. From the PowerProtect Data Manager UI left navigation pane, select **Jobs > Protection Jobs** or **Jobs > System Jobs**. The **Protection Jobs** or **Systems Jobs** window appears, displaying all jobs.

2. Select a job from the list, and then click **Export Log**.

 indicates the log export operation is in progress, and is shown next to the asset or task in the **Download Log** column. Hover over the icon 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 > Protection Jobs** or **Jobs > System Jobs**. The **Protection Jobs** or **Systems Jobs** window appears, displaying all jobs.

2. Click the job ID next to the name of the job. The **Job ID Summary** window opens.

3. Select the asset or task from the list, and then click **Export Log**.

 indicates the log export operation is in progress, and is shown next to the asset or task in the **Download Log** column. Hover over the icon to display the progress. When the log export is complete, you can download the log.

4. Click  in the **Download Log** column to download the exported log.

Modifying the System Settings

Topics:

- System settings
- System Support
- Modifying the PowerProtect Data Manager virtual machine disk settings
- Memory optimization
- Configure the DD system
- Virtual networks (VLANs)

System settings

You can use the PowerProtect Data Manager UI to modify system settings that are typically configured during PowerProtect Data Manager installation.

To access **System Settings**, click the  icon in the top-right.


Modify the network settings

Perform the following steps if you want to change the IP address of the PowerProtect Data Manager appliance, or modify other network settings such as the hostname, subnet mask, gateway, or DNS servers.

Prerequisites

 **NOTE:** When you change the IP address or hostname, the system becomes unavailable until all components are restarted.

Steps

1. From the PowerProtect Data Manager UI, click , select **System**, and then click **Network**.
2. Update the fields as necessary:
 - **Hostname**
 - **IP Address**
 - **Subnet Mask**
 - **Gateway**
 - **Primary DNS**
 - **Secondary DNS**
3. Click **Save**.

Synchronize time on 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 the systems with which it interfaces or compliance check will fail. Dell EMC recommends that all systems be configured to use an NTP server.

- NOTE:** Times in the UI are always displayed as local to the users time zone based on their browser or system settings. The PowerProtect Data Manager system might be in a different time zone but when viewing the UI it will always show the times local to the user.

Modify the appliance time zone

Use this procedure to modify the time zone for the PowerProtect Data Manager appliance.

Steps

1. From the PowerProtect Data Manager UI, click , select **System**, and then click **Timezone**.
2. From the **Timezone** list, select the applicable time zone.
3. Click **Save**.

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 to ensure successful replication.

For example, if you enable in-flight encryption in PowerProtect Data Manager, the setting must be enabled on each source and destination server before defining the PowerProtect Data Manager replication objective. If encryption is enabled after the initial definition of replication objectives, any replication jobs that were initiated during the period when the source and destination server 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 status of the in-flight encryption setting for all attached storage systems.

- NOTE:** For 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 system servers 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*.

Backup and restore encryption

You can encrypt backup or restore data that is in transit for centralized and self-service operations with DD Boost encryption, using TLS. Encryption of backup and restore data in-flight is available for application assets and NAS 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.

The following table lists the workloads and operations that support encryption of data in-flight:

NOTE: Refer to the agent user guides for more information about the centralized and self-service operations that are supported.

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)
Microsoft SQL with Application Direct	Yes	Yes (database-level restore only)	Yes	Yes (database-level restore only)
Microsoft Exchange with Application Direct	Yes	N/A	Yes	Yes
Oracle with Application Direct	Yes	N/A	Yes	Yes
SAP HANA with Application Direct	Yes	N/A	Yes	Yes
Network attached storage (NAS)	Yes	Yes	N/A	N/A

Enabling encryption imposes additional overhead. Backup and restore performance for any client could be affected by 5-20% with encryption enabled.

You can enable or disable backup and restore encryption in the PowerProtect Data Manager UI.

PowerProtect Data Manager supports backup and restore encryption for all supported DD Boost and DDOS versions. The most up-to-date software compatibility information for PowerProtect Data Manager is provided in the eLab 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 system, transmitted in encrypted form, and then decrypted before it is saved on the destination storage.

Prerequisites


Review the information in Backup and restore encryption on page 154 to learn more about backup and restore encryption.

The encryption settings determine if the data transfer is encrypted while in-flight during backup and restore operations.

- For SQL, Exchange, File System, SAP HANA, and Oracle workloads, backup and restore encryption is only supported for Application Direct hosts.
- When a new host is added to PowerProtect Data Manager, host configuration is run to push the encryption settings to the host.
- Only hosts that have PowerProtect Data Manager 19.9 application agents installed support the host configuration.

About this task

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 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 settings to the application agents on the Application Direct hosts and network attached storage (NAS).

You can disable encryption for backup and restore content by clicking the **Backup/Restore Encryption** switch. PowerProtect Data Manager creates a system job in the **Jobs > System Job** window to disable protection encryption.

Additional considerations

Review the following additional considerations for backup and restore encryption.

To validate whether encryption is being used, you can check the status of existing connections on the DD system by running the `ddboost show connections` command in the DD Boost CLI:

- The value in the **Encrypted** column is set to **Yes** if a connection has been established with encryption.
- If a client establishes a connection with encryption, and establishes another connection without encryption, the value in the **Encrypted** column is set to **Mixed**. This 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. Refer to the individual user guides for more information.
- If encryption settings exist on the DD and are also enabled in PowerProtect Data Manager, the higher encryption setting takes precedence. As a result, the **Encrypted** column will always show **Mixed** or **Yes**.

PowerProtect Data Manager licensing

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 Licensing 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 39. License types

License type	Description
Trial	Applied automatically on installation of PowerProtect Data Manager and enables full use of the product without applying 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 apply a permanent license.
Front-end protected capacity by terabyte (FETB)	The primary model of eLicensing, 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 in perpetual and term-based licenses. Your quote identifies whether your license rights are perpetual or term-based.

A perpetual license enables you to use the software for as long as 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, as long as you are in compliance with the terms of the license agreement. At the end of the license term, you must either stop using the software, extend the license term, or purchase new licenses through an agreement with Dell EMC.

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 Dell EMC license management website, you must have the License Authorization Code (LAC), which is emailed from Dell EMC. If you have not received the LAC, contact your technical support professional.

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 **Save**.

Results

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

Specify a vCenter Server as the PowerProtect Data Manager host

PowerProtect Data Manager provides an option to identify a vCenter Server as the host vCenter.

About this task


When a vCenter Server is marked as the vCenter that hosts PowerProtect Data Manager, you can use this vCenter for the following operations:

- Performing system activities, such as virtual machine-level configuration.
- Performing software updates in circumstances that require taking a PowerProtect Data Manager snapshot.
- Enabling Cloud Disaster Recovery (CDR), in order to increase the PowerProtect Data Manager CPU and memory that is required for these operations. A vCenter host is a prerequisite for CDR, as specified in the **Cloud Disaster Recovery** tab of the PowerProtect Data Manager UI **Infrastructure > Asset Sources** window.

To specify a vCenter Server as the vCenter that hosts PowerProtect Data Manager, you can:

- Add and discover this vCenter as an asset source in the PowerProtect Data Manager UI **Infrastructure > Asset Sources** window, or
- Enter the vCenter Server information in the **Hosting vCenter** window, as outlined in the following procedure.


Steps

1. From the PowerProtect Data Manager UI, 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, 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, a  icon displays next to this vCenter in the **Infrastructure > Asset Sources** window.

System Support

You can use the PowerProtect Data Manager UI to manage and modify support settings, such as the mail server setup and Secure Remote Services registration, that are typically configured during installation.

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 Technical 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 Technical Support.

SupportAssist provides the following features and benefits:

- Proactive monitoring and issue prevention
- Facilitates update package downloads
- Automatic support case creation based on event alerting
- Automatic health checks
- Communicates telemetry data
- Real-time troubleshooting
- Customer support

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

Migrating to SupportAssist

SupportAssist provides automated support capabilities for PowerProtect Data Manager systems. SupportAssist replaces Secure Remote Services (SRS) in this release of PowerProtect Data Manager.

If you have configured SRS previously, the PowerProtect Data Manager system automatically migrates SRS to SupportAssist when you update PowerProtect Data Manager.

If you do not have SRS configured, you can configure SupportAssist directly.

Use the following procedures to configure SupportAssist.

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 the SupportAssist Enterprise

Establish a connection to the SupportAssist Enterprise to ensure access to Technical Support. SupportAssist enables you to connect PowerProtect Data Manager directly or through a gateway server.

Prerequisites

- Apply a valid PowerProtect Data Manager license.
- If you are connecting through the gateway server, the SRS gateway version must be 3.40 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, and then enter the SupportAssist Access Key and PIN.
 **NOTE:** Remote Support functionality is currently not supported for PowerProtect Data Manager systems using a direct connection.
 - **Connect via Gateway**
Select this option to connect PowerProtect Data Manager through a gateway server, and then perform the following tasks.
 - a. Enter the SupportAssist gateway server IP address and port number.
 - b. Click **Test** to test the connection to the gateway server.

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.
 - c. Enter the SupportAssist Access Key and PIN.
4. Click **Enable Connect**.

Results

PowerProtect Data Manager is connected to the SupportAssist Enterprise.

Update or configure contact data

Provide contact information for the person that Technical 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.

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.

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 Dell EMC Support. Generate SupportAssist access key and PIN on page 159 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.

Test gateway connection failed

If the Secure Remote Services (SRS) gateway version is earlier than 3.40, the connection to the gateway might fail. If you are using a gateway version that is earlier than 3.40, the SRS gateway configuration is not transferred to SupportAssist after updating PowerProtect Data Manager.

When updating PowerProtect Data Manager, the precheck dialog box displays a warning indicating that the SRS gateway version in use is earlier than 3.40, and to update the SRS gateway to the compatible version.

If you are using a gateway version earlier than 3.40, the update fails with the following error:

SYS0034

Unable to upgrade from Secure Remote Services to SupportAssist.


Details

The upgrade to SupportAssist is unsuccessful for one or more of the following reasons: 1) The SupportAssist service cannot start. 2) The gateway is not accessible. 3) An issue occurs during Gen3 key upgrade.

Recommended Action

In the PowerProtect Data Manager UI: 1) To open the Support dialog, click **Settings** and select **Support**. 2) In the left pane, select **SupportAssist** to set up SupportAssist.

If this issue occurs, perform the following:

1. Check that the gateway version is 3.40 or later.
2. Set up SupportAssist. In the PowerProtect Data Manager UI, click  select **Support**, and then click **SupportAssist**.

Connection status changes to "Not Connected"

If the connection status changes to "Not Connected":

1. Ensure that all prerequisites are met in [Connect to the SupportAssist Enterprise](#) on page 159.
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 Technical 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:

- Cloud Data Recovery
- Asset Sources
- Hosts Information
- DD Inventory
- PowerProtect Data Manager operational inventory
- Integrated Storage
- Usage
- Licensing
- Compliance in last 24 hours
- Traffic Metrics
- Protection Policies
- Alerts
- Cloud Disaster Recovery metrics
- Service Level Agreement
- Assets
- Storage Systems
- Data targets
- Protection Details
- Compliance Details
- Audit logs
- Report Generated Time
- Update Summaries

CloudIQ reporting

When you enable AutoSupport, you also enable reporting. CloudIQ is a no-cost SaaS/cloud-based management application that proactively monitors and measures the overall health of Dell EMC 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 the **Reporting** menu item, or 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 UI, 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.
 - f. **Password:**
The password associated with the PowerProtect Data Manager SMTP email server.
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 [Automatically check for an update package on page 210](#)

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 UI, click  select **Support**, 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 **Alerts > 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**.
The **Alerts** window displays alert information in a table. You can filter the alerts by Severity, Date, Category, or Acknowledge.
2. Select the **System** tab.
The **System** tab displays all alert types.
3. To view more details about a specific entry, click  next to the entry in the table.
4. 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.
5. To acknowledge the alert, select the alerts and then click **Acknowledge**.
6. To add or edit a note for the alert, click **Add/Edit Note**, and when finished, click **Save**.
7. To export a report of alert information to a .CSV file which you can download for Excel, select an entry in the table and then click **Export**.

 **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 **Note** field, and click **Save**.
 4. Click **Export**.

Monitor system state and system health

In addition to the summary system health view provided in the PowerProtect Data Manager UI's **Dashboard** window, the **System Settings > Support** window provides a further breakdown of PowerProtect Data Manager system health information.

Monitor system component health

Through the **Settings** window, you can monitor the state of the appliance and the health of each system component.

To view the health of system components, click , select **Support**, and then click **System Health**.

The following table provides a summary of each component state:

Table 40. Component status

Status	Description
Running	This state appears when the associated service or component is running with full functionality. When all components are in running state, the state of the appliance is operational.
Initializing	This state appears when the component is starting. When the component 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 states. When other components are in maintenance, the appliance state is also maintenance.
Quiesce	This state appears when the service that is 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.

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 installation 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 with certificates from an approved certificate authority. This guide also contains instructions for establishing certificate-based trust with external components and systems.

Modifying the PowerProtect Data Manager virtual machine disk settings

Follow the steps in this section, under the guidance and recommendations of Dell EMC Support, to expand the size of the data disk and system disk, and modify the memory configuration.

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

- Perform the following steps from the **vSphere Web Client**:
 - Right-click the VM Direct appliance and select **Shut Down Guest OS**.
 - After the power off completes, right-click the appliance and select **Edit Settings**. The **Edit Settings** window appears with the **Virtual Hardware** button selected.
 - 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.
 - Right-click the VM Direct appliance and select **Power On**.
- 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.
 - Reboot the appliance by typing `reboot`.
 - On the **GNU GRUB** menu, press `Esc` to edit the GNU GRUB menu.
 - 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

- Press **Ctrl-x** to reboot into single-user mode.
- When prompted, type the password for the root account.
- 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 *sah* 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 *sah* 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**.

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 minimum amount of memory to assign to the PowerProtect Data Manager virtual machine in a standard environment.

Table 41. Minimum memory requirements

Deployment Type	Memory
Default	20 GB
With the Cloud Disaster Recovery (Cloud DR) Add-On	24 GB

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 32 GB of RAM. This amount of memory should not be changed after it is deployed.

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 memory

Adjust the amount of memory assigned to the PowerProtect Data Manager 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**.

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.

Supported scenarios

PowerProtect Data Manager supports virtual networks for the following use cases:

- Virtual machine backups
- Database backups
- Exchange 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 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.
- Configure the network switch port for trunk mode. This setting allows the port to carry traffic for multiple VLANs.
- Enable Virtual Guest Tagging (VGT) 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. For standard port groups, configure the virtual switch port for VLAN ID 4095, which makes all VLANs accessible. Alternatively, you can use VLAN trunking, which supports specifying multiple VLANs by ID or range. The VMware ESXi documentation provides more information.

- 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.

Dell Technologies recommends 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, Dell Technologies strongly recommends that you use the same network name in both locations for each virtual network. Record each network name for later use.

Dell Technologies also recommends 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.

Name the virtual network for protection storage

After you add protection storage, name the virtual network between the PowerProtect Data Manager and the protection storage system. To rename a virtual network (edit the network name), repeat these steps.

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 > Name Network**.
The **Name Network** window opens and displays a list of known network interfaces, assigned IP addresses, and link speeds.
3. Identify the interfaces for each new virtual network, and then type names for the virtual networks in the corresponding fields.
4. Click **Save**.
The PowerProtect Data Manager stores the network names.

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 a 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. In the **Network Name** field, type the name of the new virtual network.
Dell Technologies recommends that you keep the network names consistent for each VLAN.
4. In the **VLAN ID** field, type the numeric value 1 through 4094 that corresponds to the VLAN which this virtual network represents.
5. Provide the applicable **Subnet Mask** and **MTU** (maximum transmission unit) values for the virtual network.
Allowable **MTU** values range from 1500 to 9000.
6. For the **Static IP Pools** field, provide the indicated number of reserved IP addresses for PowerProtect Data Manager to use for communication on this virtual network.

To add values to the pool, type an IP address or range, and then click **Save**. To remove values from the pool, select a value from the pool, and then click **Delete**.

You can type each IP address separately, or you can provide an IP address range in the form **10.1.1.4-10.1.1.10**.

7. Verify that the static IP address pool contains enough addresses to add the virtual network.
8. Click **Next**.
The **Add Network** wizard moves to the **Routes** page.
9. 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.
 - 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 routes.
 - 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.
10. Click **Next**.
The **Add Network** wizard moves to the **Summary** page.
11. 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.

i 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 or does not contain the VLAN ID, you can view the details to further identify the virtual network before making changes.

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.
3. Click  for that row.
The **Details** pane opens to the right.
This pane contains information about the virtual network configuration, such as the assigned IP address for the PowerProtect Data Manager backup interface to that network, and any configured routes.
4. Click **X** to close the details pane.

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.

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** or **Routes** sections.
The **Edit Network** wizard moves to the **Configuration** or **Routes** page.
5. Modify the appropriate network parameters, and then click **Next**.
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.

Delete a virtual network

Although optional, Dell Technologies recommends that you delete virtual networks when they are no longer required.

Prerequisites

Unassign the virtual network from any applicable assets. Disable all VM Direct Engines that are configured to use the virtual network. Disable any 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.

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.

Dell Technologies recommends 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.

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.

Protecting Virtual Machines using the Transparent Snapshot Data Mover

Topics:

- Overview of transparent snapshots for virtual machine protection
- VIB installation monitoring and management
- Transparent snapshot data mover system requirements
- Prerequisites to virtual machine protection with the Transparent Snapshot Data Mover
- Virtual machine transparent snapshot unsupported features and limitations
- Transparent Snapshot Performance and Scalability

Overview of transparent snapshots for virtual machine protection

The transparent snapshot data mover (TSDM) is a new protection mechanism in PowerProtect Data Manager 19.9 and later designed to replace the VMware vStorage API for Data Protection (VADP) protection mechanism for crash-consistent virtual machine protection.

The advantages of using the TSDM protection mechanism for virtual machine data protection include the following:

- Eliminates the latency and performance impact on the production virtual machine during the protection policy life cycle.
- Reduces the CPU, storage, and memory consumption required for backups. After the initial full backup, only incremental backups using the immediate previous snapshot will be performed.
- An external VM Direct engine is not required. The VM Direct engine embedded with PowerProtect Data Manager is sufficient.
- Automatic scaling.

VIB installation monitoring and management

The vSphere Installation Bundle (VIB) is a software package that is bundled with the PowerProtect Data Manager OVA and update package and is installed automatically on a vSphere ESXi host during the PowerProtect Data Manager 19.9 installation or update. The VIB is required to enable the transparent snapshot data mover (TSDM) for virtual machines.

An entry for the job **Performing Host Configuration (vib_install)** appears in the PowerProtect Data Manager UI during the VIB installation. During the installation, vCenter and ESXi host information is detected to verify the supported versions are installed.

You can use the **Transparent Snapshot Data Movers** tab in the **Protection Engines** window of the PowerProtect Data Manager UI to monitor and manage the installation of the VIB. This window provides a vCenter hierarchy view based on the asset sources enabled in PowerProtect Data Manager. If an ESXi host is not eligible or available for the VIB installation, the status displays as **Not Eligible** in the **Protection Engines** window.

During the creation of a crash-consistent virtual machine protection policy, the VIB is deployed automatically on the vSphere cluster being protected. If all requirements are met, TSDM is used as the default protection mechanism instead of VADP. Existing policies created in PowerProtect Data Manager 19.8 and earlier can be migrated to use TSDM, provided that the virtual machine crash-consistent policies are configured with the following options:

- **Performance optimization mode**.
- **Exclude swap files from backup** is off.
- **Enable guest file system quiescing** is off.

You can use the PowerProtect Data Manager UI to apply TSDM as the data mover for virtual machine assets.

Transparent snapshot data mover system requirements

The following software is required to automatically enable use of the Transparent Snapshot Data Mover (TSDM) for virtual machine data protection operations.

NOTE: TSDM for virtual machine protection also requires that the protection policy is a performance optimized crash-consistent policy, with the quiescing and swap file exclusion options disabled.

Table 42. Software requirements

Software required	Version supported	Notes
vCenter Server	7.0 U3	vCenter and ESXi 7.0 U3 is the minimum version that is required to use TSDM, and is scheduled for release shortly after the availability of PowerProtect Data Manager 19.9. Until this version is available and installed in the environment, TSDM is not used for virtual machine protection policies.
ESXi Server	7.0 U3	
PowerProtect Data Manager software	19.9 and later	

Prerequisites to virtual machine protection with the Transparent Snapshot Data Mover

Review the following recommendations for use of the Transparent Snapshot Data Mover (TSDM) protection mechanism for virtual machine protection.

Additional privileges required for a dedicated vCenter user account to use Transparent Snapshot Data Mover

You can use the **vSphere Client** to specify the required privileges for the dedicated vCenter user account, or you can use the **PowerCLI**, which is an interface for managing vSphere.

The following table includes the additional privileges required to use the Transparent Snapshot Data Mover (TSDM) for virtual machine protection operations.

NOTE: For the remaining privileges required for the dedicated vCenter user account, see Specify the required privileges for a dedicated vCenter user account on page 64.

Table 43. Minimum required vCenter user account privileges

Setting	vCenter 7.0.3 and later required privileges	PowerCLI equivalent required privileges
Datastore	<ul style="list-style-type: none"> Datastore > File Management 	<pre> \$privileges = @('Host.Config.Patch', 'Host.Config.Image', 'Host.Config.NetService', 'Datastore.FileManagement', 'vSphereDataProtection.Protection', 'vSphereDataProtection.Recovery', 'System.Read') </pre>
Host	<ul style="list-style-type: none"> Configuration > Patch Configuration > Image Configuration > Net Service Datastore > File Management 	
System	<ul style="list-style-type: none"> System > Read 	
vSphere Data Protection	<ul style="list-style-type: none"> Protection Recovery 	

Table 43. Minimum required vCenter user account privileges (continued)

Setting	vCenter 7.0.3 and later required privileges	PowerCLI equivalent required privileges
		<code>New-VIRole -Name 'PowerProtect' -Privilege (Get-VIPrivilege -Id \$privileges)</code>

Creating VMkernel ports

For backup and restore of virtual assets from the ESXi hosts and their respective virtual machines using the Transparent Snapshot Data Mover (TSDM) protection engine, Dell Technologies strongly recommends that you create a dedicated VMkernel port for all ESXi hosts in the cluster to facilitate data transfer.

Before you begin

- For optimal data transfer between ESXi hosts and protection storage, use the same network subnet that is used for backup storage.
- For each ESXi host in the cluster, it is recommended to use a 10G physical network adapter port for TSDM backup traffic.
- Plan a unique network subnet to use exclusively for TSDM protection engine that does not overlap with any other existing network subnets. This subnet must contain the following:
 - An IP address for each VMkernel port in each ESXi host.
 - An IP address for each port in protection storage target interfaces.

Create a VMkernel port for a standard vSwitch configuration

For each ESXi host in the cluster:

1. In the **vSphere Client**, navigate to the ESXi host and select the host.
2. Right-click the host and select **Add Networking**.
3. Select **VMkernel Network Adapter**, and then click **Next**.
4. Create a new switch, or choose an existing one, following the recommendations above. When creating a new switch, assign the NIC adapter to **Active Adapters**.
5. In the **Port Properties** settings **IP settings**, select **IPv4**, and clear all other check boxes under **Available services**.
6. In the IPv4 settings, specify VMkernel IPv4 settings following the recommendations above.

Create a VMkernel port for a Distributed vSwitch configuration

1. On the **vSphere Client** home page, click **Networking**, and then navigate to and select a distributed port group.
2. From the **Actions** menu, select **Add VMkernel Adapters**.
3. On the **Select hosts** page, click **Attached hosts**, select from the hosts that are associated with the distributed switch, and then click **OK**.
4. Click **Next**.
5. On the **Configure VMkernel adapter** page, select **IPv4**, and clear all other check boxes under **Available services**.
6. In the IPv4 settings, specify VMkernel IPv4 settings following the recommendations above.

Virtual machine transparent snapshot unsupported features and limitations

Review the following unsupported features and limitations for the transparent snapshot data mover (TSDM) in PowerProtect Data Manager 19.9.

Unsupported virtual machine configurations

The following virtual machines and configurations are not supported for TSDM virtual machine protection:

- vVOL Datastores
- Physical RDMs
- Virtual RDMs
- Encrypted virtual machines
- Fault Tolerant virtual machines.

Virtual Machine Disk (VMDK) limit for virtual machines protected with TSDM

TSDM-based protection supports a maximum of 40 VMDKs per virtual machine. If this limit is exceeded, backups will be queued for a longer period of time, and will have to be cancelled manually.

For virtual machines with more than 40 VMDKs, you can override the protection mechanism at the asset level to use VADP. The section *Migrating assets to use the Transparent Snapshot Data Mover* on page 71 provides more information.

vMotion of TSDM protected virtual machines

vSphere disables the vMotion migration of virtual machines to an ESXi host version previous to 7.0 U3 when the virtual machine is protected with TSDM. In order to migrate the TSDM protected virtual machine to an ESXi version that does not support TSDM, you must disable the Lightweight Delta (LWD) filter that is attached to the virtual machine during the initial protection policy configuration. To disable the filter, remove the virtual machine from the TSDM protected virtual machine protection policy. Once the virtual machine is removed from the policy, a job is automatically initiated to disable the filter.

Once the vMotion completes, you can re-add the virtual machine to the protection policy. This virtual machine will then be protected by the VADP protection mechanism, since the new ESXi/cluster host version is lower than the version required by TSDM.

Removal of managed snapshots required prior to running virtual machine protection policies

A PowerProtect Data Manager virtual machine protection policy cannot be configured to use the TSDM protection mechanism when the virtual machine contains managed snapshots. Verify that no managed snapshots exist for the virtual machine, and then retry the configuration job from the **System Jobs** window of the PowerProtect Data Manager UI.

TSDM only available for virtual machine crash-consistent policies

Use of the TSDM protection mechanism is currently only supported for crash-consistent virtual machine protection policies. Also, the virtual machine crash-consistent policy must use the **Performance** optimization mode, with swap file exclusion and quiescing turned off.

Cloud Disaster Recovery not supported

Cloud Disaster Recovery (CDR) is not supported for TSDM virtual machine backups up PowerProtect Data Manager 19.9.

Transparent Snapshot Performance and Scalability

Review the following information related to performance considerations to scale your environment.

- ① **NOTE:** As a VMware infrastructure best practice, Dell Technologies recommends spreading the workload across ESXi Servers as much as possible. With the Transparent Snapshot Data Mover protection mechanism, you can move backup data in streams from multiple ESXi Servers.

Table 44. Scalability limits for the vCenter and ESXi Server

Component	Maximum limit
Number of protected virtual machines per ESXi Server	Unlimited
Number of VMDKs that can be protected per ESXi Server	1000
Size of VMDK	64 TB
Transparent Snapshot Data Mover (TSDM)	Up to 3000 virtual machine backups, and up to 180 concurrent virtual machine backups. ① NOTE: An external VM Direct engine is not required when using TSDM as the protection mechanism for crash-consistent virtual machine protection. For application consistent and application aware virtual machine protection, add an external VM Direct engine.

PowerProtect Functionality Within the vSphere Client

Topics:

- PowerProtect functionality within the vSphere Client
- Overview of the PowerProtect plug-in for the vSphere Client
- Overview of VASA and VMware Storage Policy Based Management

PowerProtect functionality within the vSphere Client

The **vSphere Client** integrates with PowerProtect Data Manager to provide the following functionality:

- **PowerProtect** portlet—When adding a vCenter Server as an asset source in the PowerProtect Data Manager UI, if you enable the **vSphere Plugin** option, a pane for **PowerProtect** appears in the **vSphere Client**. This pane provides a subset of PowerProtect Data Manager functionality, including the availability to perform a manual backup, image-level restore and file-level restore of PowerProtect Data Manager virtual machine protection policies.
- Storage policy association with a PowerProtect Data Manager virtual machine protection policy—vSphere Storage APIs for Storage Awareness (VASA) leverages VMware Storage Policy Based Management (SPBM) to support data protection operations, allowing you to pair SPBM policies that are created in the **vSphere Client** with protection policies that are created in PowerProtect Data Manager. This association allows you to manage all virtual machine storage and protection requirements in a centralized location (the **vSphere Client**), instead of requiring multiple user interfaces.

Overview of the PowerProtect plug-in for the vSphere Client

When adding a vCenter Server in the PowerProtect Data Manager UI, if you enable the **vSphere Plugin** option, a subset of the UI functionality becomes available within the **vSphere Client**.

The PowerProtect Data Manager portlet appears when you select **Hosts and Clusters** or **VMs and Templates** in the left pane of the **vSphere Client** home page, and then select a virtual machine within the datacenter.

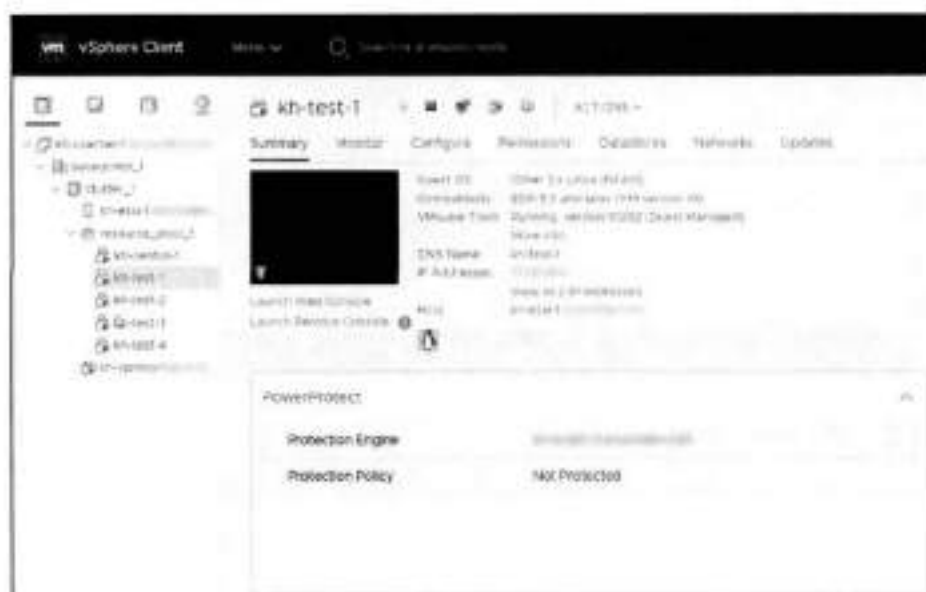


Figure 11. PowerProtect portlet in the vSphere Client

NOTE: If you were already logged into the **vSphere Client** when the vCenter discovery was started in PowerProtect Data Manager, you must log out and log back in to see the PowerProtect Data Manager UI.

If the virtual assets in the vCenter have not yet been assigned to a PowerProtect Data Manager protection policy, only the **PowerProtect** name displays in the portlet. Adding the virtual machine to a protection policy provides additional information, as shown in the following figure:



Figure 12. PowerProtect portlet with protected virtual machine

After you set up a virtual machine protection policy, you can perform the following PowerProtect Data Manager functionality within the **vSphere Client**:

- View information about protection policies and information about available protection copies.
- Monitor in-progress backup and restore operations for the virtual machine protection policy. You can also view information for successfully completed protection copies that are available for restore.
- Perform a manual backup.
- Perform an image-level restore (Restore to Original, Restore to New, or Instant Access).
- Perform a file-level restore.

Prerequisites for enabling the vSphere Client PowerProtect plug-in

To use the **vSphere Client PowerProtect** plug-in for backup and restore operations, complete the following tasks in the **vSphere Client** and the PowerProtect Data Manager UI.

- Add the vCenter Server—in the PowerProtect Data Manager UI, select **Infrastructure > Asset Sources**, and move the **vSphere Plugin** slider to the right to enable the plug-in. Add a VMware vCenter Server on page 62 provides information.

- Add privileges for the **Virtual machine power user** group (if you are already an administrator, this task is optional)—in the **vSphere Client**, go to **Administration > Roles**, select the **Virtual Machine power user (PPDM)**, and then open the **Edit Role** window.

Add the following PowerProtect Data Manager privileges:

- o Backup
- o File Level Restore to Original
- o Instant Access
- o Restore to New
- o Restore to Original



Figure 13. PowerProtect privileges added for the Virtual machine power user

NOTE: If you edit the vCenter Server in the PowerProtect Data Manager UI to unregister the **vSphere Plugin** for PowerProtect Data Manager, these PowerProtect Data Manager privileges are not removed from the user group.

- For the virtual asset (virtual machine, cluster, host) and all its child elements, add permissions to the **Virtual machine power user** group that you enabled with PowerProtect Data Manager privileges. To add these permissions, select the asset in the left pane of the **vSphere Client**, and then click the **Permissions** tab.
- Add a virtual machine protection policy in the PowerProtect Data Manager UI **Protection > Protection Policies** window to schedule a backup of the virtual machines. Add a protection policy for virtual-machine protection on page 77 provides information.

Monitor PowerProtect Data Manager virtual machine protection copies

You can use the **Monitor** tab in the **vSphere Client** to view PowerProtect Data Manager protection copies that are available for restore, and monitor in-progress backup and restore operations for the PowerProtect Data Manager virtual machine protection policy.

With a virtual machine selected, in the **Monitor** tab's navigation pane, select **PowerProtect > Protection Copies** to view information about completed PowerProtect Data Manager protection policy backups for this virtual machine. This view is the same as the view in the PowerProtect Data Manager UI **Infrastructure** window. A copy map enables you to view the available protection copies when you click on the storage icon, as described in More options for managing virtual-machine backups on page 85.

To view the status of active backup and restore operations initiated from the PowerProtect Data Manager UI or the **vSphere Client**, click the arrows icon in the lower right corner of the window to expand the **Recent Tasks** pane. You can also view this pane from the **Summary** window.

Manual PowerProtect policy backup in the vSphere Client

You can back up one or more PowerProtect Data Manager virtual machine protection policies at any time by performing a manual backup in the **vSphere Client**.

Prerequisites

- Ensure that you are logged in to the **vSphere Client** as an administrator.
- Add the **Backup** privilege to the **Administrator** group in the **vSphere Client**. To add the **Backup** privilege, complete the following steps:
 1. Select **Administration > Roles**.
 2. Select **Administrator**, and then click **Privileges** in the right pane.
 3. In the **PowerProtect Backup** section, select **Backup**.
- Ensure that virtual machine assets have been added to a virtual machine protection policy. You cannot perform manual backups of unprotected virtual machines.

Steps

1. In the left pane of the **vSphere Client** home page, select **Hosts and Clusters** or **VMs and Templates**, and then select a virtual machine within the datacenter.
The **Summary** window displays.
2. Perform a manual backup of a virtual machine protection policy by using one of the following methods:
 - In the left pane, right-click the virtual machine, and then select **PowerProtect > Backup**.
 - Within the **PowerProtect** portlet, click **Backup Now**.The **vSphere Client** starts the backup operation. A message appears indicating whether the request was processed successfully.

Results

An entry for the backup job appears in the **Jobs > Protection** window of the PowerProtect Data Manager UI. To view the status of operations, you can also click the arrows icon in the lower right corner of the window to expand the **Recent Tasks** pane.

Image-level restore of a PowerProtect backup in the vSphere Client

You can use the **vSphere Client PowerProtect** plug-in to perform an image-level restore of a PowerProtect Data Manager virtual machine protection policy backup.

About this task

Available image-level restore options in the **vSphere Client** include:

- **Restore to Original**—Restore the virtual machine to the original location on the same vCenter.
- **Restore Individual Virtual Disks**—Restore selected VMDKs to the original location on the same vCenter.
- **Restore to New**—Restore the virtual machine to a new location on the original vCenter.
- **Instant Access**—Restore the backup as a live virtual machine to view the backup and then determine whether you want to do a full restore. Instant Access sessions are made available for a default period of 7 days, which can be extended.

Steps

1. In the left pane of the **vSphere Client** home page, select **Hosts and Clusters** or **VMs and Templates**, and then select a virtual machine within the datacenter.
2. In the **Summary** window, access the backup copy by using one of the following methods:
 - In the left pane, right-click the virtual machine, and then select **PowerProtect > Restore**.
 - Within the **PowerProtect** portlet, click **Restore**.
3. On the **Select Copy** page, for each virtual machine that is listed in the table, select the radio button next to the virtual machine and click **Choose Copy**.
The **Choose Copy** dialog appears.

i **NOTE:** If you click **Next** without choosing a copy, the most recent backup copy is used.

4. In the **Choose Copy** dialog:
 - a. Select the storage icon to access the backup copies.
 - b. Choose from one of the available copies that appears in the table.
 - c. Click **OK** to close the dialog and return to the **Select Copy** page.
 - d. Click **Next**.
5. On the **Purpose** page, select from one of the following options:
 - Restore Entire VMs—Select this option if you want to restore the entire virtual machine.
 - Restore Individual Virtual Disks—Select this option if you want to restore only specific virtual machine disks (VMDKs).

i **NOTE:** Individual VMDKs can only be restored to the original location.

6. Click **Next**.
If restoring entire virtual machines, the **Restore Type** page appears. If restoring individual VMDKs, the **Select Disks** page appears.
7. On the **Restore Type** page, select from one of the available restore types.
 - For Instant Access restore, review the section Instant access virtual machine restore on page 117.
 - For Restore to New, review the section Restore to a new virtual machine on page 115.
 - For Restore to Original, review the section Restore to the original virtual machine on page 112.
 - For Restore Individual Virtual Disks, review the section Restore individual virtual disks on page 114.

The wizard updates to display the options specific to the restore type that you selected.

i **NOTE:** Options such as vCenter, resource pool, and datastore are limited to the logged-in vSphere user's permissions, and are not necessarily the same as a PowerProtect Data Manager administrator.

8. Click **Next**. The **Summary** page appears.
9. Review your selections and then click **Restore**.

Results

An entry for the restore job appears in the **Recent Tasks** pane of the **vSphere Client** and in the **Restore > Running Sessions** window of the PowerProtect Data Manager UI.

Next steps

For Instant Access restores, when the virtual machine is powered on and you select the virtual machine in the left pane of the **Summary** window, the session information appears within the **PowerProtect** portlet. If you need extra time for this session, you can click **Extend Session** and increase session availability by up to 7 days.

File-level restore of a PowerProtect backup in the vSphere Client

You can use the **PowerProtect** portlet in the **vSphere Client** to perform a file-level restore of a PowerProtect Data Manager virtual machine protection policy backup.

Prerequisites

Note the following before performing file-level restore in the **vSphere Client**:

- A minimum vCenter version 6.7 U1 is required.
- Review the section Supported platform versions for file-level restore for supported platform and operating system versions.
- Review the section File-level restore and SQL restore limitations on page 247.
- Ensure that the **FLR Agent** is installed on the target virtual machine by logging into the virtual machine and verifying that the agent package is installed and the agent process is running. If the **FLR Agent** is not installed, the installation is initiated automatically when you start the mount.

When installing the **FLR Agent** on Windows virtual machines, the user must be an administrator account. When installing the **FLR Agent** on Linux virtual machines, the user must be the root user account. The section **FLR Agent for virtual machine file level restore** on page 248 provides more information.

i **NOTE:**
For file-level restores, you can only restore files:

- From a Windows backup to a Windows machine, or from a Linux backup to a Linux machine.
- To virtual machines within the same vCenter.

About this task

Available file-level restore options in the **vSphere Client** include:

- Restore single or multiple files to the original folder and overwrite the original files within the same virtual machine, or
- Restore single or multiple files to a new folder with a new name within the same virtual machine.

Steps

1. In the left pane of the **vSphere Client** home page, select **Hosts and Clusters** or **VMs and Templates**, and then select a virtual machine within the datacenter.
The **Summary** window displays.

2. Access the backup copy by using one of the following methods:

- In the left pane, right-click the virtual machine, and then select **PowerProtect > File Level Restore**.
- Within the **PowerProtect** portlet, click **File Level Restore**.

3. On the **Select Copy** page, for each virtual machine that is listed in the table, select the radio button next to the virtual machine and click **Choose Copy**.

The **Choose Copy** dialog appears.

NOTE: If you click **Next** without choosing a copy, the most recent backup copy is used.

4. In the **Choose Copy** dialog:

- a. Select the storage icon to access the backup copies.
- b. Choose from one of the available copies that appears in the table:
- c. Click **OK** to close the dialog and return to the **Select Copy** page.
- d. Click **Next**.

5. On the **Mount Copy** page:

- a. To initiate the disk mount, type the guest operating system user credentials:
 - If there are administrator-level credentials associated with the virtual assets or protection policy being restored, specify end-user credentials.
 - If there are no administrator-level credentials associated with the virtual assets or protection policy being restored, specify administrator credentials. These credentials will be handled as end-user credentials.
- b. (Optional) Leave **Keep FLR Agent Installed** selected when you want the FLR Agent to remain on the destination virtual machine after the restore completes.
- c. Click **Start Mount** to initiate the disk mount.

If not already installed, the **FLR Agent** is installed on the target virtual machine. A progress bar indicates when the mount completes.

NOTE: You cannot browse the contents of the virtual machine backup until the mounting of the destination virtual machine completes successfully.

- d. Upon successful mount, click **Next**.

6. On the **Select Files to Recover** page:

- a. Expand individual folders to browse the original virtual machine backup, and select the objects that you want to restore to the destination virtual machine.
- b. Click **Next**.

NOTE: In the browse view, each directory or hard drive appears twice. Selecting an object from one location selects the object in the duplicate location as well.

7. On the **Options** page, select from one of the following options:

- Restore to Original Folder and Overwrite Original Files—Select this option to restore all selected files to their original location on the original virtual machine.
- Restore to an Alternate Folder—Select this option if you want to restore to a new folder in a new location on the original virtual machine.

8. Click **Next**.

If performing the restore to the original virtual machine, the **Summary** page displays. You can go to the final step. If performing the restore to an alternate location on the original virtual machine, the **Restore Location** page displays.

9. On the **Restore Location** page:
 - a. Browse the folder structure of the virtual machine to select the new folder where you want to restore the objects.
 - b. Click **Next**.
10. On the **Summary** page:
 - a. Review the information to ensure that the restore details are correct. You can click **Edit** next to the **Restore Location** or **Files Selected** rows to change the information.
 - b. Click **Restore**.

Results

An entry for the restore job appears in the **Recent Tasks** pane of the **vSphere Client** and in the **Restore > Running Sessions** window of the PowerProtect Data Manager UI.

Overview of VASA and VMware Storage Policy Based Management

vSphere Storage APIs for Storage Awareness (VASA) is a set of application program interfaces (APIs) that allow arrays to integrate with vCenter for management functionality. Storage Vendor Providers allow the vCenter Server to retrieve information from storage arrays, including topology, capabilities (such as native thin provisioning and deduplication), and status. The policy-based management functionality of a VASA provider helps administrators choose the appropriate storage device, and monitors and reports information about existing storage policies.

Starting in vSphere version 7.0 U1, VASA support is extended to Data Protection operations by leveraging VMware Storage Policy Based Management (SPBM). SPBM spans all storage offerings from VMware, allowing policies to provision and manage storage for any virtual machine application. The integration of PowerProtect Data Manager and SPBM allows you to:

- Pair SPBM policies with protection policies, allowing you to meet virtual machine storage and protection requirements within vSphere without requiring the PowerProtect Data Manager UI for data protection operations.
- Add new or existing virtual assets to an SPBM policy. You can also reassign these assets and remove them from the policy.
- View policy compliance status, including data protection policy information.
- Protect virtual machines at scale, allowing you to manage capacity resources and overcome challenges such as capacity planning and different service level requirements.

Enabling VASA and SPBM within the **vSphere Client** for integration with PowerProtect Data Manager requires you to perform the following:

- Register the VASA provider to allow for storage provisioning information flow between PowerProtect Data Manager and the vCenter Server.
- Select the PowerProtect Data Manager storage awareness provider within the vCenter Server storage policy component creation workflow, which exposes the list of available PowerProtect Data Manager virtual machine protection policies.
- Assign the PowerProtect Data Manager protection policy to an SPBM policy, which is automatically assigned to virtual machines when they are represented by an instance.
- Monitor the status of storage compliance of the virtual assets protected by these PowerProtect Data Manager policies.

If you replace the default self-signed security certificates for PowerProtect Data Manager with certificates from an approved certificate authority, you must exchange the new security certificates with vCenter. The *PowerProtect Data Manager Security Configuration Guide* provides instructions.

Register the VASA provider for policy association

The following procedure describes how to register the VASA provider to enable PowerProtect Data Manager communication with the vCenter Server and use the provider to enable an association between a virtual machine storage policy and a PowerProtect Data Manager virtual machine protection policy.

Prerequisites

The vSphere version must be a minimum 7.0 U1.

Steps

1. In the **vSphere Client**, go to **Menu > Hosts and Clusters**.

1. In the left pane, select the vCenter Server, and then select the **Configure** tab.
3. Under **Security**, select **Storage Providers**, and then click **+ Add**.
The **New Storage Provider** dialog appears.
4. On the **New Storage Provider** dialog:
 - a. Specify a name for the provider.
 - b. Specify a URL in the format `https://my-ppdm.example.com:9009/vasa/version.xml`, where `my-ppdm.example.com` is the PowerProtect Data Manager fully qualified hostname.
 - c. Provide PowerProtect Data Manager credentials for a user with the Administrator role, and then click **OK**.
These credentials are only required for the initial login to perform the registration. Subsequent log-in attempts use certificates.

If the vCenter Server does not trust the SSL certificate of the PowerProtect Data Manager server, a prompt appears, asking if you want to accept the certificate as trusted. You can trust this certificate, or alternatively, you can securely obtain a copy of the certificate as a file, and then click **Browse** within this prompt to select and trust the certificate. The vCenter documentation provides more information.

NOTE: For self-signed or untrusted certificates, an error might appear. You can dismiss and ignore this error.

5. Provide PowerProtect Data Manager administrator level credentials, and then click **OK**.
The dialog updates to indicate that the registration is in progress. If the vCenter Server does not trust the SSL certificate of the PowerProtect Data Manager server, a prompt displays to accept the certificate as trusted. You can trust this certificate, or alternatively, you can securely obtain a copy of the certificate as a file, and then click **Browse** within this prompt to select and trust the certificate. The vCenter documentation provides more information.

NOTE: For self-signed or untrusted certificates, an error might appear. You can ignore this error.
6. When the registration is complete, click **OK** to exit the **New Storage Provider** dialog.
The **Configure** tab updates to display the new VASA provider.

Results

You can now use the **vSphere Client** to create a virtual machine storage policy and associate this policy with an existing PowerProtect Data Manager virtual machine protection policy.

NOTE: If the provider goes offline at any point, you can select the provider in the table and click **Rescan** to reestablish a connection. Also, if the provider is removed and then readded, any policies that were previously assigned to the provider are restored.

Add an SPBM policy and associate with a PowerProtect Data Manager virtual machine policy

Use the **vSphere Client** to create a virtual machine storage policy and associate this policy with an existing PowerProtect Data Manager virtual machine protection policy.

Steps

1. In the **vSphere Client**, select the vCenter Server in the left pane.
2. Go to **Menu > Policies and Profiles**.
3. In the left pane, select **VM Storage Policies**, and then click **Create** in the right pane.
The **Create VM Storage Policy** wizard appears.
4. Provide a name and description that helps identify this policy as a storage policy that you want to associate with a PowerProtect Data Manager protection policy, and then click **Next**.
5. On the **Policy Structure** page, select **Enable host based rules**, and then click **Next**.
6. On the **Host based services** page, select the **Data Protection** tab, and then perform the following:
 - a. Select **Custom**.
 - b. From the **Provider** list, select **DellEMC PowerProtect** as the registered provider.
 - c. From the **PPDM Protection Policy** list, select an existing PowerProtect Data Manager virtual machine protection policy that you want to associate with this storage policy.

NOTE: Dell Technologies recommends that you use a descriptive name for the PowerProtect Data Manager virtual machine protection policy so that the purpose is easy to identify, since the **vSphere Client** does not provide policy

details within the **PowerProtect** portlet. If you decide to rename the PowerProtect Data Manager policy at any point, the association is retained since the UUID of the policy is used to create the connection.

d. Click **Next**.

7. Complete the storage policy details, and click **Finish**.

Results

The **VM Storage Policies** window displays the new storage policy in the table. An association is created between the PowerProtect Data Manager policy and the virtual machine storage policy, and the **PowerProtect** portlet in the **vSphere Client** updates to display the PowerProtect Data Manager protection policy. You can now perform manual backups and scheduled restores of the virtual assets in this policy.

When you assign the new storage policy to a virtual machine, that virtual machine should automatically be assigned to the associated PowerProtect Data Manager protection policy as well. Also, if you are creating a new virtual machine, you can assign a storage policy to the new virtual machine during this process.

NOTE: You can create separate storage policies for each virtual machine disk, but only the policy that is associated with the virtual machine is used for data protection.

NOTE: If you want to remove a virtual machine from protection, assign the virtual machine to a different policy, or to the Detastore Default policy.

Monitor virtual machine protection policy compliance

You can use the **Storage Policies** portlet within the **vSphere Client** to monitor the compliance of virtual assets in PowerProtect Data Manager virtual machine protection policies.

To access the portlet:

- Select the **Summary** tab, or
- Select the **Configure** tab, select a virtual machine in the left pane, and then click Policies.

If a virtual asset was unassigned from the policy within PowerProtect Data Manager, the policy displays as **Non-compliant**.

VMware Cloud (VMC) on Amazon Web Services (AWS)

Topics:

- PowerProtect Data Manager image backup and recovery
- Supported PowerProtect Data Manager and DDVE deployment configurations
- Deployment and configuration best practices and requirements
- Configuring the VMC-on-AWS portal
- Interoperability with PowerProtect Data Manager features
- vCenter server inventory requirements
- Creating a dedicated cloud-based vCenter user account
- Add a VM Direct Engine
- Unsupported operations

PowerProtect Data Manager image backup and recovery

PowerProtect Data Manager provides image backup and restore support for VMware Cloud (VMC) on Amazon Web Services (AWS).

Using PowerProtect Data Manager to protect virtual-machine assets in VMC on AWS is similar to how you protect virtual-machine assets in an on-premises data center. The following sections provide information on network configuration requirements, PowerProtect Data Manager best practices, and unsupported PowerProtect Data Manager operations.

Supported PowerProtect Data Manager and DDVE deployment configurations

In order to protect virtual-machine assets in VMC on AWS, PowerProtect Data Manager and DDVE can be deployed in several ways.

When deploying PowerProtect Data Manager and DDVE, two possible deployment environments are VMware Cloud on AWS (VMC on AWS) and the AWS Marketplace (AWS). The following table describes the supported deployment configurations of the two products:

Table 45. Supported deployment configurations

PowerProtect Data Manager	DDVE
VMware Cloud on AWS	VMware Cloud on AWS
VMware Cloud on AWS	AWS Marketplace
AWS Marketplace	AWS Marketplace

When deploying PowerProtect Data Manager to VMC on AWS, an Open Virtualization Appliance (OVA) is used. This puts PowerProtect Data Manager into the VMC-on-AWS environment in order to protect the VMware assets. When deploying PowerProtect Data Manager to AWS, a machine image is used. This puts PowerProtect Data Manager into a cloud-marketplace environment, but still allows the VMware assets in the VMC-on-AWS environment to be protected.

For more information about the different deployment types, see the *PowerProtect Data Manager Deployment Guide* and the *PowerProtect Data Manager AWS Deployment Guide*.

Deployment and configuration best practices and requirements

Deploying and configuring PowerProtect Data Manager, DDVE, and other components in a certain way provides an efficient protection of virtual-machine assets.

To perform data protection and disaster recovery tasks in VMC on AWS, consider the following recommendations for the backup infrastructure:

- Deploy PowerProtect Data Manager and DDVE either to VMC on AWS or to AWS.
- Deploy the VM Direct appliance to VMC on AWS.
- Deploy at least one VM Direct appliance for each software-defined data center (SDDC) cluster in the VMC-on-AWS environment.
- When deploying or configuring PowerProtect Data Manager or the VM Direct appliance, ensure that the DNS server IP points to the internal DNS server that is running in vCenter inventory.
- Ensure that the internal DNS server has both forward and reverse lookup entries for all of the required components, such as the PowerProtect Data Manager server, the VM Direct appliance, and the DDVE appliance.
- If using NSX-T, add the vCenter server to PowerProtect Data Manager by using the FQDN.
- If using NSX-V, add the vCenter server to PowerProtect Data Manager by using the public FQDN of the vCenter server.
- When adding the vCenter server to PowerProtect Data Manager, perform one of the following actions:
 - Specify the login credentials for the `cloudadmin@vmc.local` user.
 - Refer to [Creating a dedicated cloud-based vCenter user account](#) on page 193 to create a dedicated cloud-based vCenter user account, and then specify the login credentials for that user.
- You can clone backups to another instance of DDVE running in the same environment as the first instance. This type of deployment enables backup copies to be stored for longer retention, leveraging the AWS network for transferring data at lower latency and cost when compared to the public Internet.
- You can store backups outside of the VMC-on-AWS environment. For example, store backups on an AWS virtual private cloud (VPC). This type of deployment enables efficient data transfer over the fast ENI connection that is used by VMware to communicate with AWS.

Configuring the VMC-on-AWS portal

Domain Name System (DNS) resolution is critical for deployment and configuration of PowerProtect Data Manager, the PowerProtect Data Manager external proxy, and DDVE. All infrastructure components should be resolvable through a fully qualified domain name (FQDN). Resolvable means that components are accessible through both forward (A) and reverse (PTR) lookups.

Ensure that the VMC-on-AWS portal meets the following requirements:

- By default, there is no external access to the vCenter server in the software-defined data center (SDDC). You can open access to the vCenter server by configuring a firewall rule. To enable communication to the vCenter public IP address from the SDDC logical network, set the firewall rule in the compute gateway of VMC on AWS. If the firewall rule is not configured in the SDDC, PowerProtect Data Manager does not allow you to add the vCenter server.
- The default compute gateway firewall rules prevent all virtual machine traffic from reaching the Internet. To enable the PowerProtect Data Manager virtual machine to connect to the Internet, create a compute gateway firewall rule. This action enables outbound traffic on the logical network to which the PowerProtect Data Manager server virtual machine is connected.
- Configure DNS to allow machines in the SDDC to resolve FQDNs to their public IP addresses. If the DNS server is not configured in the SDDC, the PowerProtect Data Manager server does not allow you to add the vCenter server by using the server's public FQDN or IP address.
- It is recommended that you deploy the DD system as a virtual appliance. If deploying DDVE to VMC-on-AWS, connect the SDDC to an AWS account during the SDDC creation, and then select a VPC and subnet within that account.
- DDVE must be connected to the SDDC through the VMC-on-AWS Elastic Network Interfaces (ENIs). This action allows the SDDC, the services in the VPC, and subnet in the AWS account to communicate without having to route traffic through the Internet gateway.
- The same ENI channel is recommended for access to DDVE.

For more information about configuring ENIs, see <https://vmc.vmware.com/console/aws-link>.

- If DDVE is running in VMC-on-AWS, configure the inbound and outbound firewall rules of the compute gateway for DDVE connectivity.

For detailed information on what incoming and outgoing ports need to be opened for PowerProtect-VM proxy solution, refer to the *PowerProtect Data Manager Security Configuration Guide*.

- If using NSX-T, configure DNS to resolve to the internal IP address of the vCenter server. Navigate to **SDDC Management > Settings > vCenter FQDN**, and then select the **Private vCenter IP address** to directly access the management network over the built-in firewall.
- Open TCP port 443 of the vCenter and ESXi servers in both the management and compute gateways.
- If DDVE is running in VMC-on-AWS, the inbound and outbound firewall rules of the VMC-on-AWE VPC security group are configured to provide connectivity between the SDDC compute gateway and DDVE.
- If there is replication between DDVE instances, ensure the following:
 - The security group in AWS is configured to allow all inbound traffic from the private IPs of the DDVE instances
 - The DDVE instances can ping each other using their FQDNs

Interoperability with PowerProtect Data Manager features

VMC on AWS has certain restrictions on workloads and resource pools. To ensure proper operation, select the **Workload and Compute** sections in AWS.

Do not use the following non-accessible areas:

- vSANdatastore datastore
- Management VMs folder in VMs and Templates view
- Mgmt-ResourcePool resource pool in Hosts and Clusters view

vCenter server inventory requirements

In the vCenter server inventory of the SDDC, ensure that the following requirements are met:

- An internal DNS name server must be running inside vCenter inventory. This will be referenced by all the workloads running in the SDDC.
- The internal DNS server must have **Forwarders** enabled to access the internet. This action is required to resolve the vCenter server's public FQDN. Forwarders are DNS servers that the server can use to resolve DNS queries for records that the server itself cannot resolve.

Creating a dedicated cloud-based vCenter user account

It is recommended that you set up a separate vCenter user account at the root level of the vCenter hierarchy. This account is strictly dedicated for use with PowerProtect Data Manager and the VM Direct protection engine in cloud-based environments.

Use of a generic user account such as **Administrator** could make future troubleshooting efforts difficult as it might not be clear which **Administrator** actions are actually interfacing or communicating with PowerProtect Data Manager. Using a separate vCenter user account ensures maximum clarity if it becomes necessary to examine vCenter logs.

You can specify the credentials for a vCenter user account when you add the vCenter server as an asset source in the user interface. When you add the vCenter server, ensure that you specify a user whose cloud-based role is defined at the vCenter level and not restricted to a lower-level container object in the vSphere object hierarchy.

Specify the required privileges for a dedicated cloud-based vCenter user account

You can use the **vSphere Client** to specify the required privileges for the dedicated cloud-based vCenter user account, or you can use the **PowerCLI**, which is an interface for managing vSphere.

The following table includes the privileges required for this user.

NOTE: For the privileges required when administering on-premises PowerProtect Data Manager, see Specify the required privileges for a dedicated vCenter user account on page 64.

Table 46. Minimum required cloud-based vCenter user account privileges

Setting	vCenter 6.0 and later required privileges	PowerCLI equivalent required privileges
Alarms	<ul style="list-style-type: none"> Create alarm Modify alarm 	<pre>\$privileges = @('System.Anonymous', 'System.View', 'System.Read', 'Alarm.Create', 'Alarm.Edit', 'Cryptographer.Access', 'Datastore.Browse', 'Datastore.DeleteFile', 'Datastore.FileManagement', 'Datastore.AllocateSpace', 'Datastore.Config', 'Folder.Create', 'Global.ManageCustomFields', 'Global.SetCustomField', 'Global.LogEvent', 'Global.CancelTask', 'InventoryService.Tagging.AttachTag', 'InventoryService.Tagging.ObjectAttachable', 'InventoryService.Tagging.CreateTag', 'InventoryService.Tagging.CreateCategory', 'Network.Assign', 'Resource.AssignVMToPool', 'Resource.HotMigrate', 'Resource.ColdMigrate', 'Sessions.ValidateSession', 'StorageProfile.View', 'VApp.ApplicationConfig', 'VApp.Export', 'VApp.Import', 'VirtualMachine.Config.Rename', 'VirtualMachine.Config.Annotation', 'VirtualMachine.Config.AddExistingDisk', 'VirtualMachine.Config.AddNewDisk', 'VirtualMachine.Config.RemoveDisk', 'VirtualMachine.Config.RawDevice', 'VirtualMachine.Config.HostUSBDevice', 'VirtualMachine.Config.CPUCount', 'VirtualMachine.Config.Memory', 'VirtualMachine.Config.AddRemoveDevice', 'VirtualMachine.Config.EditDevice', 'VirtualMachine.Config.Settings', 'VirtualMachine.Config.Resource', 'VirtualMachine.Config.UpgradeVirtualHardware', 'VirtualMachine.Config.ResetGuestInfo', 'VirtualMachine.Config.AdvancedConfig', 'VirtualMachine.Config.DiskLease', 'VirtualMachine.Config.SwapPlacement', 'VirtualMachine.Config.DiskExtend', 'VirtualMachine.Config.ChangeTracking', 'VirtualMachine.Config.ReloadFromPath', 'VirtualMachine.Config.ManagedBy', 'VirtualMachine.GuestOperations.Query', 'VirtualMachine.GuestOperations.Modify', 'VirtualMachine.GuestOperations.Execute', 'VirtualMachine.Interact.PowerOn', 'VirtualMachine.Interact.PowerOff',</pre>
Cryptographic operations	<ul style="list-style-type: none"> Direct Access <p>NOTE: This only applies to AVS and GCVE.</p>	
Datastore	<ul style="list-style-type: none"> Allocate space Browse datastore Configure datastore Low level file operations Remove file 	
Folder	<ul style="list-style-type: none"> Create folder 	
Global	<ul style="list-style-type: none"> Cancel task Log event Manage custom attributes Set custom attribute 	
vSphere Tagging	<ul style="list-style-type: none"> Assign or Unassign vSphere Tag Assign or Unassign vSphere Tag on Object <p>NOTE: This only applies to vCenter 7.0 and later.</p> <ul style="list-style-type: none"> Create vSphere Tag Create vSphere Tag Category 	
Network	<ul style="list-style-type: none"> Assign network 	
Resource	<ul style="list-style-type: none"> Assign virtual machine to resource pool Migrate powered off virtual machine Migrate powered on virtual machine 	
Sessions	<ul style="list-style-type: none"> Validate session 	
SPBM policy restore	<ul style="list-style-type: none"> Profile-driven storage view 	
vApp	<ul style="list-style-type: none"> Export Import vApp application configuration 	
Virtual Machine		
Change Configuration	<ul style="list-style-type: none"> Acquire disk lease Add existing disk Add new disk Add or remove device Advanced configuration Change CPU count Change Memory Change Settings Change Swapfile placement Change resource Configure Host USB device Configure Raw device Configure managedby 	

Table 46. Minimum required cloud-based vCenter user account privileges (continued)

Setting	vCenter 6.0 and later required privileges	PowerCLI equivalent required privileges
	<ul style="list-style-type: none"> Extend virtual disk Modify device settings Reload from path Remove disk Rename Reset guest information Set annotation Toggle disk change tracking Upgrade virtual machine compatibility 	<pre>'VirtualMachine.Interact.Reset', 'VirtualMachine.Interact.ConsoleInteract t', 'VirtualMachine.Interact.DeviceConnecti on', 'VirtualMachine.Interact.SetCDMedia', 'VirtualMachine.Interact.ToolsInstall', 'VirtualMachine.Interact.GuestControl', 'VirtualMachine.Inventory.Create', 'VirtualMachine.Inventory.Register', 'VirtualMachine.Inventory.Delete', 'VirtualMachine.Inventory.Unregister', 'VirtualMachine.Provisioning.DiskRandom Access', 'VirtualMachine.Provisioning.DiskRandom Read', 'VirtualMachine.Provisioning.GetVmFiles ', 'VirtualMachine.Provisioning.MarkAsTemp late', 'VirtualMachine.State.CreateSnapshot', 'VirtualMachine.State.RevertToSnapshot') 'VirtualMachine.State.RemoveSnapshot')</pre>
Edit Inventory	<ul style="list-style-type: none"> Create new Register Remove Unregister 	
Guest operations	<ul style="list-style-type: none"> Guest operation modifications Guest operation program execution Guest operation queries 	
Interaction	<ul style="list-style-type: none"> Configure CD media Connect devices Console interaction Guest operating system management by VIX API Install VMware Tools Power off Power on Reset 	<pre>New-VIRole -Name 'PowerProtect' -Privilege (Get-VIPrivilege -Id \$privileges)</pre>
Provisioning	<ul style="list-style-type: none"> Allow disk access Allow read-only disk access Allow virtual machine download Mark as template 	
Snapshot Management	<ul style="list-style-type: none"> Create snapshot Remove snapshot Revert to snapshot 	

Add a VM Direct Engine

Perform the following steps in the **Protection Engines** window of the PowerProtect Data Manager UI to deploy an external VM Direct Engine, also referred to as a VM proxy. The VM Direct Engine facilitates data movement for virtual-machine protection policies.

Prerequisites

Review the sections Requirements for an external VM Direct Engine on page 66 and Transport mode considerations on page 252.

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

About this task

The PowerProtect Data Manager software comes bundled with an embedded VM Direct engine, which is automatically used as a fallback proxy for performing backups and restores when the added external proxies fail or are disabled. It is recommended that you deploy external proxies by adding a VM Direct Engine for the following reasons:

- An external VM Direct Engine for VM proxy backup and recovery can provide improved performance and reduce network bandwidth utilization by using source-side deduplication.

- The embedded VM Direct engine has limited capacity for backup streams.
- The embedded VM Direct engine is not supported for VMC-on-AWS, AVS-on-Azure, or GCVE-on-GCP operations.

NOTE: Cloud-based deployments of PowerProtect Data Manager do not support the configuration of data-traffic routing or VLANs. Skip the **Networks Configuration** page.

Steps

1. From the left navigation pane, select **Infrastructure > Protection Engines**.
The **Protection Engines** window appears.
2. In the **VM Direct Engines** pane of the **Protection Engines** window, click **Add**.
The **Add Protection Engine** wizard displays.
3. On the **Protection Engine Configuration** page, complete the required fields, which are marked with an asterisk.
 - **Hostname, Gateway, IP Address, Netmask, and Primary DNS**—Note that only IPv4 addresses are supported.
 - **vCenter to Deploy**—If you have added multiple vCenter Server instances, select the vCenter on which to deploy the protection engine.

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 protection engine.
 - **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 based on ranking (whether the datastores are shared or local), and available capacity (the datastore with the most capacity appearing at the top of the list).

You can choose the specific datastore on which the protection engine resides, or leave the default selection of **<automatic>** to allow PowerProtect Data Manager to determine the best location to host the protection engine.
 - **Transport Mode**—Select **Hot Add**.
 - **Supported Protection Type**—Select whether this protection engine is intended for **Virtual Machine, Kubernetes Tanzu guest cluster, or NAS** asset protection.
4. Click **Next**.
5. Click **Next** to skip the **Networks Configuration** page.
6. On the **Summary** page, review the information and then click **Finish**.
The protection engine is added to the **VM Direct Engines** pane. An additional column indicates the engine purpose. Note that it can take several minutes to register the new protection engine in PowerProtect Data Manager. The protection engine also appears in the **vSphere Client**.

Results

When an external VM Direct Engine is deployed and registered, PowerProtect Data Manager uses this engine instead of the embedded VM Direct engine for any data protection operations that involve virtual machine protection policies. If every external VM Direct Engine is unavailable, PowerProtect Data Manager uses the embedded VM Direct engine as a fallback to perform limited scale backups and restores. If you do not want to use the external VM Direct Engine, you can disable this engine. Additional VM Direct actions on page 68 provides more information.

NOTE: The external VM Direct Engine is always required for VMC-on-AWS, AVS-on-Azure, and GCVE-on-GCP operations. If no external VM Direct Engine is available for these solutions, data protection operations fail.

Next steps

If the protection engine deployment fails, review the network configuration of PowerProtect Data Manager in the **System Settings** window to correct any inconsistencies in network properties. After successfully completing the network reconfiguration, delete the failed protection engine and then add the protection engine in the **Protection Engines** window.

When configuring the VM Direct Engine in a VMC-on-AWS, AVS-on-Azure, or GCVE-on-GCP environment, if you deploy the VM Direct Engine to the root of the cluster instead of inside the Compute-ResourcePool, you must move the VM Direct Engine inside the Compute-ResourcePool.

Unsupported operations

PowerProtect Data Manager image backup and restore in VMC on AWS does not currently support the following operations:

- PowerProtect Search functionality
- The vSphere Storage Policy Based Management (SPBM) integration with PowerProtect Data Manager
- A VM Direct appliance that is configured with dual-stack or IPv6
- Application-consistent data protection for Microsoft SQL with the VM Direct appliance
- VM Backup and Recovery HTML5 plug-in functionality for vSphere
- Image-based backups and restores that use NBD or the NBDSSL transport mode
- Image-based backups and restores when a datacenter is placed inside a folder in the SDDC
- File-level recoveries of an image-based backup
- Instant-access restores of an image-based backup
- Emergency restores of an image-based restore directly to an ESXi host, bypassing the vCenter server
- Backup and restore operations with anything other than the **CloudAdmin** role or a customized role that has all of the privileges listed in *Specify the required privileges for a dedicated cloud-based vCenter user account* on page 193

NOTE: If protecting virtual-machine assets with a PowerProtect Data Manager machine image deployed to AWS, Cloud Disaster Recovery (CDR) and Search Clusters are also unsupported.

Azure VMware Solution (AVS) on Microsoft Azure

Topics:

- PowerProtect Data Manager image backup and recovery
- Supported PowerProtect Data Manager and DDVE deployment configurations
- Deployment and configuration best practices and requirements
- Configuring the AVS-on-Azure portal
- vCenter server inventory requirements
- Creating a dedicated cloud-based vCenter user account
- Add a VM Direct Engine
- Unsupported operations

PowerProtect Data Manager image backup and recovery

PowerProtect Data Manager provides image backup and restore support for Azure VMware Solution (AVS) on Microsoft Azure. Using PowerProtect Data Manager to protect virtual-machine assets AVS on Azure is similar to how you protect virtual-machine assets in an on-premises data center. This section provides information on network configuration requirements, PowerProtect Data Manager best practices, and unsupported PowerProtect Data Manager operations.

Supported PowerProtect Data Manager and DDVE deployment configurations

In order to protect virtual-machine assets in AVS on Azure, PowerProtect Data Manager and DDVE can be deployed in a couple of ways.

When deploying PowerProtect Data Manager and DDVE, two possible deployment environments are Azure VMware Solution (AVS on Azure) and the Azure Marketplace (Azure). The following table describes the supported deployment configurations of the two products:

Table 47. Supported deployment configurations

PowerProtect Data Manager	DDVE
Azure VMware Solution	Azure Marketplace
Azure Marketplace	Azure Marketplace

When deploying PowerProtect Data Manager to AVS on Azure, an Open Virtualization Appliance (OVA) is used. This puts PowerProtect Data Manager into the AVS-on-Azure environment in order to protect the VMware assets. When deploying PowerProtect Data Manager to Azure, a machine image is used. This puts PowerProtect Data Manager into a cloud-marketplace environment, but still allows the VMware assets in the AVS-on-Azure environment to be protected.

For more information about the different deployment types, see the *PowerProtect Data Manager Deployment Guide* and the *PowerProtect Data Manager Azure Deployment Guide*.

Deployment and configuration best practices and requirements

Deploying and configuring PowerProtect Data Manager, DDVE, and other components in a certain way provides an efficient protection of virtual-machine assets.

To perform data protection and disaster recovery tasks in AVS on Azure, consider the following recommendations and requirements for the backup infrastructure:

- Deploy PowerProtect Data Manager either to AVS on Azure or to Azure.
- Deploy DDVE to Azure.
- Deploy the VM Direct appliance to AVS on Azure. Deploy at least one VM Direct appliance for each software-defined data center (SDDC) cluster in the AVS-on-Azure environment.
- When deploying or configuring PowerProtect Data Manager or the VM Direct appliance, ensure that the DNS server IP points to the internal DNS server that is running in vCenter inventory.
- Ensure that the internal DNS server has both forward and reverse lookup entries for all of the required components, such as the PowerProtect Data Manager server, the VM Direct appliance, and DDVE.
- If using NSX-T, add the vCenter server to PowerProtect Data Manager by using the FQDN.
- If using NSX-V, add the vCenter server to PowerProtect Data Manager by using the public FQDN of the vCenter server.
- When adding the vCenter server to PowerProtect Data Manager, perform one of the following actions:
 - Specify the login credentials for the `cloudadmin@vsphere.local` user.
 - Refer to [Creating a dedicated cloud-based vCenter user account](#) on page 193 to create a dedicated cloud-based vCenter user account, and then specify the login credentials for that user.
- You can clone backups to another instance of DDVE running in Azure. This type of deployment enables backup copies to be stored for longer retention, leveraging the Azure network for transferring data at lower latency and cost when compared to the public Internet.

Configuring the AVS-on-Azure portal

Domain Name System (DNS) resolution is critical for deployment and configuration of PowerProtect Data Manager, the PowerProtect Data Manager external proxy, and the DDVE appliance. All infrastructure components should be resolvable through a fully qualified domain name (FQDN). Resolvable means that components are accessible through both forward (A) and reverse (PTR) lookups.

Ensure that the AVS-on-Azure portal meets the following requirements:

- If you have deployed a PowerProtect Data Manager OVA to AVS on Azure or a PowerProtect Data Manager machine image to Azure, it is configured to use a custom DNS server.
 - ❗ **NOTE:** If you have already deployed PowerProtect Data Manager without a custom DNS server, you will have to redeploy it. For more information, see the [PowerProtect Data Manager Deployment Guide](#) or the [PowerProtect Data Manager Azure Deployment Guide](#).
- Forward and reverse DNS lookups exist for PowerProtect Data Manager, vCenter, DDVE, ESXi, and each VM Direct Engine.
- DNS is configured to allow machines in the SDDC to resolve FQDNs to their IP addresses.
- DDVE is running in Azure. If you have more than one DDVE instance running in Azure to perform replication, the DDVE instances have the ability to ping each other using their FQDNs.
 - ❗ **NOTE:** DDVE running in AVS-on-Azure is not supported.
- DDVE has DNS entries for PowerProtect Data Manager and each VM Direct Engine.
- SDDC is connected to an Azure account, and an Azure cloud and subnet within that account is selected.
- Any DDVE instance on Azure is connected to the SDDC through a Vnet. This action allows the SDDC, the services in the Azure cloud, and subnets in the Azure account to communicate without having to route traffic through the Internet gateway.

The same Vnets are recommended for access to DDVE instances. For more information about configuring Vnets, see [About Virtual Network](#).

vCenter server inventory requirements

In the vCenter server inventory of the SDDC, ensure that the following requirements are met:

- An internal DNS name server must be running inside vCenter inventory. This will be referenced by all the workloads running in the SDDC.
- The internal DNS server must have **Forwarders** enabled to access the internet. This action is required to resolve the vCenter server's public FQDN. Forwarders are DNS servers that the server can use to resolve DNS queries for records that the server itself cannot resolve.

Creating a dedicated cloud-based vCenter user account

It is recommended that you set up a separate vCenter user account at the root level of the vCenter hierarchy. This account is strictly dedicated for use with PowerProtect Data Manager and the VM Direct protection engine in cloud-based environments.

Use of a generic user account such as **Administrator** could make future troubleshooting efforts difficult as it might not be clear which **Administrator** actions are actually interfacing or communicating with PowerProtect Data Manager. Using a separate vCenter user account ensures maximum clarity if it becomes necessary to examine vCenter logs.

You can specify the credentials for a vCenter user account when you add the vCenter server as an asset source in the user interface. When you add the vCenter server, ensure that you specify a user whose cloud-based role is defined at the vCenter level and not restricted to a lower-level container object in the vSphere object hierarchy.

Specify the required privileges for a dedicated cloud-based vCenter user account

You can use the **vSphere Client** to specify the required privileges for the dedicated cloud-based vCenter user account, or you can use the **PowerCLI**, which is an interface for managing vSphere.

The following table includes the privileges required for this user.

NOTE: For the privileges required when administering on-premises PowerProtect Data Manager, see Specify the required privileges for a dedicated vCenter user account on page 64.

Table 48. Minimum required cloud-based vCenter user account privileges

Setting	vCenter 6.0 and later required privileges	PowerCLI equivalent required privileges
Alarms	<ul style="list-style-type: none">• Create alarm• Modify alarm	<pre>\$privileges = @('System.Anonymous', 'System.View', 'System.Read', 'Alarm.Create', 'Alarm.Edit', 'Cryptographer.Access', 'Datastore.Browse', 'Datastore.DeleteFile', 'Datastore.FileManagement', 'Datastore.AllocateSpace', 'Datastore.Config', 'Folder.Create', 'Global.ManageCustomFields', 'Global.SetCustomField', 'Global.LogEvent', 'Global.CancelTask', 'InventoryService.Tagging.AttachTag', 'InventoryService.Tagging.ObjectAttachable', 'InventoryService.Tagging.CreateTag', 'InventoryService.Tagging.CreateCategory'</pre>
Cryptographic operations	<ul style="list-style-type: none">• Direct Access <p>NOTE: This only applies to AVS and GCVE.</p>	
Datastore	<ul style="list-style-type: none">• Allocate space• Browse datastore• Configure datastore• Low level file operations• Remove file	
Folder	<ul style="list-style-type: none">• Create folder	
Global	<ul style="list-style-type: none">• Cancel task• Log event• Manage custom attributes• Set custom attribute	
vSphere Tagging	<ul style="list-style-type: none">• Assign or Unassign vSphere Tag	

Table 48. Minimum required cloud-based vCenter user account privileges (continued)

Setting	vCenter 6.0 and later required privileges	PowerCLI equivalent required privileges
	<ul style="list-style-type: none"> Assign or Unassign vSphere Tag on Object <p>NOTE: This only applies to vCenter 7.0 and later.</p> <ul style="list-style-type: none"> Create vSphere Tag Create vSphere Tag Category 	<pre>'Network.Assign', 'Resource.AssignVMToPool', 'Resource.HotMigrate', 'Resource.ColdMigrate', 'Sessions.ValidateSession', 'StorageProfile.View', 'VApp.ApplicationConfig', 'VApp.Export', 'VApp.Import', 'VirtualMachine.Config.Rename', 'VirtualMachine.Config.Annotation', 'VirtualMachine.Config.AddExistingDisk', 'VirtualMachine.Config.AddNewDisk', 'VirtualMachine.Config.RemoveDisk', 'VirtualMachine.Config.RawDevice', 'VirtualMachine.Config.HostUSBDevice', 'VirtualMachine.Config.CPUCount', 'VirtualMachine.Config.Memory', 'VirtualMachine.Config.AddRemoveDevice', 'VirtualMachine.Config.EditDevice', 'VirtualMachine.Config.Settings', 'VirtualMachine.Config.Resource', 'VirtualMachine.Config.UpgradeVirtualHardware', 'VirtualMachine.Config.ResetGuestInfo', 'VirtualMachine.Config.AdvancedConfig', 'VirtualMachine.Config.DiskLease', 'VirtualMachine.Config.SwapPlacement', 'VirtualMachine.Config.DiskExtend', 'VirtualMachine.Config.ChangeTracking', 'VirtualMachine.Config.ReloadFromPath', 'VirtualMachine.Config.ManagedBy', 'VirtualMachine.GuestOperations.Query', 'VirtualMachine.GuestOperations.Modify', 'VirtualMachine.GuestOperations.Execute', 'VirtualMachine.Interact.PowerOn', 'VirtualMachine.Interact.PowerOff', 'VirtualMachine.Interact.Reset', 'VirtualMachine.Interact.ConsoleInteraction', 'VirtualMachine.Interact.DeviceConnection', 'VirtualMachine.Interact.SetCDMedia', 'VirtualMachine.Interact.ToolsInstall', 'VirtualMachine.Interact.GuestControl', 'VirtualMachine.Inventory.Create', 'VirtualMachine.Inventory.Register', 'VirtualMachine.Inventory.Delete', 'VirtualMachine.Inventory.Unregister', 'VirtualMachine.Provisioning.DiskRandomAccess', 'VirtualMachine.Provisioning.DiskRandomRead', 'VirtualMachine.Provisioning.GetVmFiles', 'VirtualMachine.Provisioning.MarkAsTemplate', 'VirtualMachine.State.CreateSnapshot', 'VirtualMachine.State.RevertToSnapshot', 'VirtualMachine.State.RemoveSnapshot' }</pre>
Network	<ul style="list-style-type: none"> Assign network 	
Resource	<ul style="list-style-type: none"> Assign virtual machine to resource pool Migrate powered off virtual machine Migrate powered on virtual machine 	
Sessions	<ul style="list-style-type: none"> Validate session 	
SPBM policy restore	<ul style="list-style-type: none"> Profile-driven storage view 	
vApp	<ul style="list-style-type: none"> Export Import vApp application configuration 	
Virtual Machine		
Change Configuration	<ul style="list-style-type: none"> Acquire disk lease Add existing disk Add new disk Add or remove device Advanced configuration Change CPU count Change Memory Change Settings Change Swapfile placement Change resource Configure Host USB device Configure Raw device Configure managedby Extend virtual disk Modify device settings Reload from path Remove disk Rename Reset guest information Set annotation Toggle disk change tracking Upgrade virtual machine compatibility 	<pre>'VirtualMachine.Config.Resource', 'VirtualMachine.Config.UpgradeVirtualHardware', 'VirtualMachine.Config.ResetGuestInfo', 'VirtualMachine.Config.AdvancedConfig', 'VirtualMachine.Config.DiskLease', 'VirtualMachine.Config.SwapPlacement', 'VirtualMachine.Config.DiskExtend', 'VirtualMachine.Config.ChangeTracking', 'VirtualMachine.Config.ReloadFromPath', 'VirtualMachine.Config.ManagedBy', 'VirtualMachine.GuestOperations.Query', 'VirtualMachine.GuestOperations.Modify', 'VirtualMachine.GuestOperations.Execute', 'VirtualMachine.Interact.PowerOn', 'VirtualMachine.Interact.PowerOff', 'VirtualMachine.Interact.Reset', 'VirtualMachine.Interact.ConsoleInteraction', 'VirtualMachine.Interact.DeviceConnection', 'VirtualMachine.Interact.SetCDMedia', 'VirtualMachine.Interact.ToolsInstall', 'VirtualMachine.Interact.GuestControl', 'VirtualMachine.Inventory.Create', 'VirtualMachine.Inventory.Register', 'VirtualMachine.Inventory.Delete', 'VirtualMachine.Inventory.Unregister', 'VirtualMachine.Provisioning.DiskRandomAccess', 'VirtualMachine.Provisioning.DiskRandomRead', 'VirtualMachine.Provisioning.GetVmFiles', 'VirtualMachine.Provisioning.MarkAsTemplate', 'VirtualMachine.State.CreateSnapshot', 'VirtualMachine.State.RevertToSnapshot', 'VirtualMachine.State.RemoveSnapshot' }</pre>
Edit Inventory	<ul style="list-style-type: none"> Create new Register Remove Unregister 	
Guest operations	<ul style="list-style-type: none"> Guest operation modifications Guest operation program execution Guest operation queries 	
Interaction	<ul style="list-style-type: none"> Configure CD media Connect devices Console interaction 	<pre>New-VIRole -Name 'PowerProtect'</pre>

Table 48. Minimum required cloud-based vCenter user account privileges (continued)

Setting	vCenter 6.0 and later required privileges	PowerCLI equivalent required privileges
	<ul style="list-style-type: none"> • Guest operating system management by VIX API • Install VMware Tools • Power off • Power on • Reset 	<pre>-Privilege (Get-VIPrivilege -Id \$privileges)</pre>
Provisioning	<ul style="list-style-type: none"> • Allow disk access • Allow read-only disk access • Allow virtual machine download • Mark as template 	
Snapshot Management	<ul style="list-style-type: none"> • Create snapshot • Remove snapshot • Revert to snapshot 	

Add a VM Direct Engine

Perform the following steps in the **Protection Engines** window of the PowerProtect Data Manager UI to deploy an external VM Direct Engine, also referred to as a VM proxy. The VM Direct Engine facilitates data movement for virtual-machine protection policies.

Prerequisites

Review the sections **Requirements for an external VM Direct Engine** on page 66 and **Transport mode considerations** on page 252.

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

About this task

The PowerProtect Data Manager software comes bundled with an embedded VM Direct engine, which is automatically used as a fallback proxy for performing backups and restores when the added external proxies fail or are disabled. It is recommended that you deploy external proxies by adding a VM Direct Engine for the following reasons:

- An external VM Direct Engine for VM proxy backup and recovery can provide improved performance and reduce network bandwidth utilization by using source-side deduplication.
- The embedded VM Direct engine has limited capacity for backup streams.
- The embedded VM Direct engine is not supported for VMC-on-AWS, AVS-on-Azure, or GCVE-on-GCP operations.

NOTE: Cloud-based deployments of PowerProtect Data Manager do not support the configuration of data-traffic routing or VLANs. Skip the **Networks Configuration** page.

Steps

1. From the left navigation pane, select **Infrastructure > Protection Engines**.
The **Protection Engines** window appears.
2. In the **VM Direct Engines** pane of the **Protection Engines** window, click **Add**.
The **Add Protection Engine** wizard displays.
3. On the **Protection Engine Configuration** page, complete the required fields, which are marked with an asterisk:
 - **Hostname, Gateway, IP Address, Netmask, and Primary DNS**—Note that only IPv4 addresses are supported.
 - **vCenter to Deploy**—If you have added multiple vCenter Server instances, select the vCenter on which to deploy the protection engine.

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 protection engine.
 - **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 based on ranking (whether the datastores are shared or local), and available capacity (the datastore with the most capacity appearing at the top of the list).

You can choose the specific datastore on which the protection engine resides, or leave the default selection of **<automatic>** to allow PowerProtect Data Manager to determine the best location to host the protection engine.

- **Transport Mode**—Select **Hot Add**.
- **Supported Protection Type**—Select whether this protection engine is intended for **Virtual Machine**, **Kubernetes** Tanzu guest cluster, or **NAS** asset protection.

4. Click **Next**.
5. Click **Next** to skip the **Networks Configuration** page.
6. On the **Summary** page, review the information and then click **Finish**.

The protection engine is added to the **VM Direct Engines** pane. An additional column indicates the engine purpose. Note that it can take several minutes to register the new protection engine in PowerProtect Data Manager. The protection engine also appears in the **vSphere Client**.

Results

When an external VM Direct Engine is deployed and registered, PowerProtect Data Manager uses this engine instead of the embedded VM Direct engine for any data protection operations that involve virtual machine protection policies. If every external VM Direct Engine is unavailable, PowerProtect Data Manager uses the embedded VM Direct engine as a fallback to perform limited scale backups and restores. If you do not want to use the external VM Direct Engine, you can disable this engine. Additional VM Direct actions on page 68 provides more information.

- NOTE:** The external VM Direct Engine is always required for VMC-on-AWS, AVS-on-Azure, and GCVE-on-GCP operations. If no external VM Direct Engine is available for these solutions, data protection operations fail.

Next steps

If the protection engine deployment fails, review the network configuration of PowerProtect Data Manager in the **System Settings** window to correct any inconsistencies in network properties. After successfully completing the network reconfiguration, delete the failed protection engine and then add the protection engine in the **Protection Engines** window.

When configuring the VM Direct Engine in a VMC-on-AWS, AVS-on-Azure, or GCVE-on-GCP environment, if you deploy the VM Direct Engine to the root of the cluster instead of inside the Compute-ResourcePool, you must move the VM Direct Engine inside the Compute-ResourcePool.

Unsupported operations

PowerProtect Data Manager image backup and restore in AVS on Azure does not currently support the following operations:

- PowerProtect Search functionality
- A VM Direct appliance that is configured with dual-stack or IPv6
- Application-consistent data protection for Microsoft SQL with the VM Direct appliance
- VM Backup and Recovery HTML5 plug-in functionality for vSphere
- Image-based backups and restores that use NBD or the NBDSSL transport mode
- Image-based backups and restores when a datacenter is placed inside a folder in the SDDC
- File-level recoveries of an image-based backup
- Instant-access restores of an image-based backup
- Emergency restores of an image-based restore directly to an ESXi host, bypassing the vCenter server
- Backup and restore operations with anything other than the **CloudAdmin** role or a customized role that has all of the privileges listed in Specify the required privileges for a dedicated cloud-based vCenter user account on page 193

- NOTE:** If protecting virtual-machine assets with a PowerProtect Data Manager machine image deployed to Azure, Cloud Disaster Recovery (CDR), Search Clusters, and Microsoft Exchange are also unsupported.

Google Cloud VMware Engine (GCVE) on Google Cloud Product (GCP)

Topics:

- PowerProtect Data Manager image backup and recovery
- Supported PowerProtect Data Manager and DDVE deployment configurations
- Deployment and configuration best practices and requirements
- Configuring the GCVE-on-GCP portal
- vCenter server inventory requirements
- Creating a dedicated cloud-based vCenter user account
- Add a VM Direct Engine
- Unsupported operations

PowerProtect Data Manager image backup and recovery

PowerProtect Data Manager provides image backup and restore support for Google Cloud VMware Engine (GCVE) on Google Cloud Platform (GCP).

Using PowerProtect Data Manager to protect virtual-machine assets in GCVE on GCP is similar to how you protect virtual-machines assets in an on-premises data center. The following sections provide information on network configuration requirements, PowerProtect Data Manager best practices, and unsupported PowerProtect Data Manager operations.

Supported PowerProtect Data Manager and DDVE deployment configurations

In order to protect virtual-machine assets in GCVE on GCP, PowerProtect Data Manager and DDVE can be deployed in a couple of ways.

When deploying PowerProtect Data Manager and DDVE, two possible deployment environments are Google Cloud VMware Engine (GCVE on GCP) and the Google Cloud Marketplace (GCP). The following table describes the supported deployment configurations of the two products:

Table 49. Supported deployment configurations

PowerProtect Data Manager	DDVE
Google Cloud VMware Engine	Google Cloud Marketplace
Google Cloud Marketplace	Google Cloud Marketplace

When deploying PowerProtect Data Manager to GCVE on GCP, an Open Virtualization Appliance (OVA) is used. This puts PowerProtect Data Manager into the GCVE-on-GCP environment in order to protect the VMware assets. When deploying PowerProtect Data Manager to GCP, a machine image is used. This puts PowerProtect Data Manager into a cloud-marketplace environment, but still allows the VMware assets in the GCVE-on-GCP environment to be protected.

For more information about the different deployment types, see the *PowerProtect Data Manager Deployment Guide* and the *PowerProtect Data Manager GCP Deployment Guide*.

Deployment and configuration best practices and requirements

For GCVE-on-GCP support, ensure that the following requirements are met:

To perform data protection and disaster recovery tasks in GCVE on GCP, consider the following recommendations and requirements for the backup infrastructure deployment:

- Deploy PowerProtect Data Manager either to GCVE on GCP or to GCP.
- Deploy DDVE to GCP.
- Deploy the VM Direct appliance in a GCVE-on-GCP environment. Deploy at least one VM Direct appliance for each software-defined data center (SDDC) cluster in GCVE on GCP.
- When deploying or configuring PowerProtect Data Manager or the VM Direct appliance, ensure that the DNS server IP points to the internal DNS server that is running in vCenter inventory.
- Ensure that the internal DNS server has both forward and reverse lookup entries for all of the required components, such as the PowerProtect Data Manager server, the VM Direct appliance, and DDVE.
- If using NSX-T, add the vCenter server to PowerProtect Data Manager by using the FQDN.
- If using NSX-V, add the vCenter server to PowerProtect Data Manager by using the public FQDN of the vCenter server.
- When adding the vCenter server to PowerProtect Data Manager, perform one of the following actions:
 - Specify the login credentials for the `CloudOwner@gve.local` user.
 - Refer to the following section to create a dedicated cloud-based vCenter user account, and then specify the login credentials for that user.
- You can clone backups to another DDVE instance running in GCP. This type of deployment enables backup copies to be stored for longer retention, leveraging the GCP network for transferring data at lower latency and cost when compared to the public Internet.

Configuring the GCVE-on-GCP portal

Domain Name System (DNS) resolution is critical for deployment and configuration of PowerProtect Data Manager, the PowerProtect Data Manager external proxy, and DDVE. All infrastructure components should be resolvable through a fully qualified domain name (FQDN). Resolvable means that components are accessible through both forward (A) and reverse (PTR) lookups.

Ensure that the GCVE-on-GCP portal meets the following requirements:

- If you have deployed a PowerProtect Data Manager OVA to GCVE on GCP or a PowerProtect Data Manager machine image to GCP, it is configured to use a custom DNS server.
 - ① **NOTE:** If you have already deployed PowerProtect Data Manager without a custom DNS server, you will have to redeploy it. For more information, see the *PowerProtect Data Manager Deployment Guide* or the *PowerProtect Data Manager GCP Deployment Guide*.
- Forward and reverse DNS lookups exist for PowerProtect Data Manager, vCenter, DDVE, ESXi, and each VM Direct Engine.
- DNS is configured to allow machines in the SDDC to resolve FQDNs to their IP addresses.
- DDVE is running in GCP. If you have more than one DDVE instance running in GCP to perform replication, both DDVE instances have the ability to ping each other using their FQDNs.
 - ① **NOTE:** DDVE running in GCVE on GCP is not supported.
- DDVE has DNS entries for PowerProtect Data Manager and each VM Direct Engine.
- SDDC is connected to a Google account, and a Google cloud and subnet within that account is selected.
- Any DDVE instances running in GCP is connected to the SDDC through a Vnet. This action allows the SDDC, the services in GCP, and subnets in GCP to communicate without having to route traffic through the Internet gateway.

The same Vnet is recommended for access to DDVE instances. For more information about configuring Vnets, see *About Virtual Network*.

vCenter server inventory requirements

In the vCenter server inventory of the SDDC, ensure that the following requirements are met:

- An internal DNS name server must be running inside vCenter inventory. This will be referenced by all the workloads running in the SDDC.
- The internal DNS server must have **Forwarders** enabled to access the internet. This action is required to resolve the vCenter server's public FQDN. Forwarders are DNS servers that the server can use to resolve DNS queries for records that the server itself cannot resolve.

Creating a dedicated cloud-based vCenter user account

It is recommended that you set up a separate vCenter user account at the root level of the vCenter hierarchy. This account is strictly dedicated for use with PowerProtect Data Manager and the VM Direct protection engine in cloud-based environments.

Use of a generic user account such as **Administrator** could make future troubleshooting efforts difficult as it might not be clear which **Administrator** actions are actually interfacing or communicating with PowerProtect Data Manager. Using a separate vCenter user account ensures maximum clarity if it becomes necessary to examine vCenter logs.

You can specify the credentials for a vCenter user account when you add the vCenter server as an asset source in the user interface. When you add the vCenter server, ensure that you specify a user whose cloud-based role is defined at the vCenter level and not restricted to a lower-level container object in the vSphere object hierarchy.

Specify the required privileges for a dedicated cloud-based vCenter user account

You can use the **vSphere Client** to specify the required privileges for the dedicated cloud-based vCenter user account, or you can use the **PowerCLI**, which is an interface for managing vSphere.

The following table includes the privileges required for this user.

NOTE: For the privileges required when administering on-premises PowerProtect Data Manager, see Specify the required privileges for a dedicated vCenter user account on page 64.

Table 50. Minimum required cloud-based vCenter user account privileges

Setting	vCenter 6.0 and later required privileges	PowerCLI equivalent required privileges
Alarms	<ul style="list-style-type: none">• Create alarm• Modify alarm	<pre>\$privileges = @('System.Anonymous', 'System.View', 'System.Read', 'Alarm.Create', 'Alarm.Edit', 'Cryptographer.Access', 'Datastore.Browse', 'Datastore.DeleteFile', 'Datastore.FileManagement', 'Datastore.AllocateSpace', 'Datastore.Config', 'Folder.Create', 'Global.ManageCustomFields', 'Global.SetCustomField', 'Global.LogEvent', 'Global.CancelTask', 'InventoryService.Tagging.AttachTag', 'InventoryService.Tagging.ObjectAttachable', 'InventoryService.Tagging.CreateTag', 'InventoryService.Tagging.CreateCategory'</pre>
Cryptographic operations	<ul style="list-style-type: none">• Direct Access <p>NOTE: This only applies to AVS and GCVE.</p>	
Datastore	<ul style="list-style-type: none">• Allocate space• Browse datastore• Configure datastore• Low level file operations• Remove file	
Folder	<ul style="list-style-type: none">• Create folder	
Global	<ul style="list-style-type: none">• Cancel task• Log event• Manage custom attributes• Set custom attribute	
vSphere Tagging	<ul style="list-style-type: none">• Assign or Unassign vSphere Tag	

Table 50. Minimum required cloud-based vCenter user account privileges (continued)

Setting	vCenter 6.0 and later required privileges	PowerCLI equivalent required privileges
	<ul style="list-style-type: none"> Assign or Unassign vSphere Tag on Object <p>① NOTE: This only applies to vCenter 7.0 and later.</p> <ul style="list-style-type: none"> Create vSphere Tag Create vSphere Tag Category 	'Network.Assign', 'Resource.AssignVMToPool', 'Resource.HotMigrate', 'Resource.ColdMigrate', 'Sessions.ValidateSession', 'StorageProfile.View', 'VApp.ApplicationConfig', 'VApp.Export', 'VApp.Import', 'VirtualMachine.Config.Rename', 'VirtualMachine.Config.Annotation', 'VirtualMachine.Config.AddExistingDisk'
Network	<ul style="list-style-type: none"> Assign network 	'VirtualMachine.Config.AddNewDisk', 'VirtualMachine.Config.RemoveDisk', 'VirtualMachine.Config.RawDevice', 'VirtualMachine.Config.HostUSBDevice', 'VirtualMachine.Config.CPUCount', 'VirtualMachine.Config.Memory', 'VirtualMachine.Config.AddRemoveDevice'
Resource	<ul style="list-style-type: none"> Assign virtual machine to resource pool Migrate powered off virtual machine Migrate powered on virtual machine 	'VirtualMachine.Config.EditDevice', 'VirtualMachine.Config.Settings', 'VirtualMachine.Config.Resource', 'VirtualMachine.Config.UpgradeVirtualHardware', 'VirtualMachine.Config.ResetGuestInfo', 'VirtualMachine.Config.AdvancedConfig', 'VirtualMachine.Config.DiskLease', 'VirtualMachine.Config.SwapPlacement', 'VirtualMachine.Config.DiskExtend', 'VirtualMachine.Config.ChangeTracking', 'VirtualMachine.Config.ReloadFromPath', 'VirtualMachine.Config.ManagedBy', 'VirtualMachine.GuestOperations.Query', 'VirtualMachine.GuestOperations.Modify'
Sessions	<ul style="list-style-type: none"> Validate session 	'VirtualMachine.GuestOperations.Execute'
SPBM policy restore	<ul style="list-style-type: none"> Profile-driven storage view 	'VirtualMachine.Interact.PowerOn', 'VirtualMachine.Interact.PowerOff', 'VirtualMachine.Interact.Reset', 'VirtualMachine.Interact.ConsoleInteract', 'VirtualMachine.Interact.DeviceConnection', 'VirtualMachine.Interact.SetCDMedia', 'VirtualMachine.Interact.ToolsInstall', 'VirtualMachine.Interact.GuestControl', 'VirtualMachine.Inventory.Create', 'VirtualMachine.Inventory.Register', 'VirtualMachine.Inventory.Delete', 'VirtualMachine.Inventory.Unregister', 'VirtualMachine.Provisioning.DiskRandomAccess', 'VirtualMachine.Provisioning.DiskRandomRead', 'VirtualMachine.Provisioning.GetVmFiles', 'VirtualMachine.Provisioning.MarkAsTemplate', 'VirtualMachine.State.CreateSnapshot', 'VirtualMachine.State.RevertToSnapshot'
vApp	<ul style="list-style-type: none"> Export Import vApp application configuration 	'VirtualMachine.State.RemoveSnapshot') New-VIRole -Name 'PowerProtect'
Virtual Machine		
Change Configuration	<ul style="list-style-type: none"> Acquire disk lease Add existing disk Add new disk Add or remove device Advanced configuration Change CPU count Change Memory Change Settings Change Swapfile placement Change resource Configure Host USB device Configure Raw device Configure managedby Extend virtual disk Modify device settings Reload from path Remove disk Rename Reset guest information Set annotation Toggle disk change tracking Upgrade virtual machine compatibility 	
Edit Inventory	<ul style="list-style-type: none"> Create new Register Remove Unregister 	
Guest operations	<ul style="list-style-type: none"> Guest operation modifications Guest operation program execution Guest operation queries 	
Interaction	<ul style="list-style-type: none"> Configure CD media Connect devices Console interaction 	

Table 50. Minimum required cloud-based vCenter user account privileges (continued)

Setting	vCenter 6.0 and later required privileges	PowerCLI equivalent required privileges
	<ul style="list-style-type: none"> • Guest operating system management by VIX API • Install VMware Tools • Power off • Power on • Reset 	<pre>-Privilege (Get-VIPrivilege -Id \$privileges)</pre>
Provisioning	<ul style="list-style-type: none"> • Allow disk access • Allow read-only disk access • Allow virtual machine download • Mark as template 	
Snapshot Management	<ul style="list-style-type: none"> • Create snapshot • Remove snapshot • Revert to snapshot 	

Add a VM Direct Engine

Perform the following steps in the **Protection Engines** window of the PowerProtect Data Manager UI to deploy an external VM Direct Engine, also referred to as a VM proxy. The VM Direct Engine facilitates data movement for virtual-machine protection policies.

Prerequisites

Review the sections *Requirements for an external VM Direct Engine* on page 66 and *Transport mode considerations* on page 252.

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

About this task

The PowerProtect Data Manager software comes bundled with an embedded VM Direct engine, which is automatically used as a fallback proxy for performing backups and restores when the added external proxies fail or are disabled. It is recommended that you deploy external proxies by adding a VM Direct Engine for the following reasons:

- An external VM Direct Engine for VM proxy backup and recovery can provide improved performance and reduce network bandwidth utilization by using source-side deduplication.
- The embedded VM Direct engine has limited capacity for backup streams.
- The embedded VM Direct engine is not supported for VMC-on-AWS, AVS-on-Azure, or GCVE-on-GCP operations.

NOTE: Cloud-based deployments of PowerProtect Data Manager do not support the configuration of data-traffic routing or VLANs. Skip the **Networks Configuration** page.

Steps

1. From the left navigation pane, select **Infrastructure > Protection Engines**.
The **Protection Engines** window appears.
2. In the **VM Direct Engines** pane of the **Protection Engines** window, click **Add**.
The **Add Protection Engine** wizard displays.
3. On the **Protection Engine Configuration** page, complete the required fields, which are marked with an asterisk.
 - **Hostname, Gateway, IP Address, Netmask, and Primary DNS**—Note that only IPv4 addresses are supported.
 - **vCenter to Deploy**—If you have added multiple vCenter Server instances, select the vCenter on which to deploy the protection engine.

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 protection engine.
 - **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 based on ranking (whether the datastores are shared or local), and available capacity (the datastore with the most capacity appearing at the top of the list).

You can choose the specific datastore on which the protection engine resides, or leave the default selection of **<automatic>** to allow PowerProtect Data Manager to determine the best location to host the protection engine.

- **Transport Mode**—Select **Hot Add**.
- **Supported Protection Type**—Select whether this protection engine is intended for **Virtual Machine, Kubernetes** Tanzu guest cluster, or **NAS** asset protection.

4. Click **Next**.
5. Click **Next** to skip the **Networks Configuration** page.
6. On the **Summary** page, review the information and then click **Finish**.

The protection engine is added to the **VM Direct Engines** pane. An additional column indicates the engine purpose. Note that it can take several minutes to register the new protection engine in PowerProtect Data Manager. The protection engine also appears in the **vSphere Client**.

Results

When an external VM Direct Engine is deployed and registered, PowerProtect Data Manager uses this engine instead of the embedded VM Direct engine for any data protection operations that involve virtual machine protection policies. If every external VM Direct Engine is unavailable, PowerProtect Data Manager uses the embedded VM Direct engine as a fallback to perform limited scale backups and restores. If you do not want to use the external VM Direct Engine, you can disable this engine. Additional VM Direct actions on page 68 provides more information.

- ① **NOTE:** The external VM Direct Engine is always required for VMC-on-AWS, AVS-on-Azure, and GCVE-on-GCP operations. If no external VM Direct Engine is available for these solutions, data protection operations fail.

Next steps

If the protection engine deployment fails, review the network configuration of PowerProtect Data Manager in the **System Settings** window to correct any inconsistencies in network properties. After successfully completing the network reconfiguration, delete the failed protection engine and then add the protection engine in the **Protection Engines** window.

When configuring the VM Direct Engine in a VMC-on-AWS, AVS-on-Azure, or GCVE-on-GCP environment, if you deploy the VM Direct Engine to the root of the cluster instead of inside the Compute-ResourcePool, you must move the VM Direct Engine inside the Compute-ResourcePool.

Unsupported operations

PowerProtect Data Manager image backup and restore in GCVE on GCP does not currently support the following operations:

- PowerProtect Search functionality
 - A VM Direct appliance that is configured with dual-stack or IPv6
 - Application-consistent data protection for Microsoft SQL with the VM Direct appliance
 - VM Backup and Recovery HTML5 plug-in functionality for vSphere
 - Image-based backups and restores that use NBD or the NBDSSL transport mode
 - Image-based backups and restores when a datacenter is placed inside a folder in the SDDC
 - File-level recoveries of an image-based backup
 - Instant-access restores of an image-based backup
 - Emergency restores of an image-based restore directly to an ESXi host, bypassing the vCenter server
 - Backup and restore operations with anything other than the **CloudOwner** role or a customized role that has all of the privileges listed in Specify the required privileges for a dedicated cloud-based vCenter user account on page 193
- ① **NOTE:** If protecting virtual-machine assets with a PowerProtect Data Manager machine image deployed to GCP, Cloud Disaster Recovery (CDR), Search Clusters, Microsoft Exchange, and block-based backups (BBB) with the File System agent (FSA) are also unsupported.

Performing Updates

Topics:

- Managing update packages
- Updating the version of PowerProtect Data Manager
- Update PowerProtect Data Manager from version 19.8 to version 19.9
- Update PowerProtect Data Manager from version 19.7 to version 19.9
- Update PowerProtect Data Manager from versions 19.3–19.6 to version 19.9

Managing update packages

You manage update packages from the **Software Update** window.

The **Software Update** window allows you to perform the following operations:

- Automatically or manually check for an update package
- Download an update package
- Upload an update package
- Delete an update package
- Perform a precheck on an update package
- Install an update package

Only the Administrator role can manage and perform updates.

Automatically check for an update package

You can configure PowerProtect Data Manager to automatically check for a new update package once a day. If a new update package is available, you can either be alerted or have it automatically downloaded.

Prerequisites

SupportAssist must be enabled. For more information, see *Configuring SupportAssist for PowerProtect Data Manager* on page 158 and *Enable or disable SupportAssist* on page 161.

Steps

1. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.
2. Click  and then select **Software Update**.
3. In the **Software Update** window, click **Check for Updates**.
4. In the **Software Update > Check for Updates** pane, click .
5. In the **Configure Updates** window, choose what happens if a new package update is available.
 - To only be notified of the package update, select **Check and notify**.
 - To automatically download the package update, select **Check and automatically download**.
6. Click **Save**.

Troubleshooting automatic downloads

Even if PowerProtect Data Manager has been configured to automatically download any new update package, there are a couple of conditions that disable this feature.

The following table describes the common issues with automatic downloads and how to resolve them.

Table 51. Common issues with automatic downloads

Issue	Reason	Resolution
Automatic downloads are disabled.	There is already an update package ready to install and displayed in the Software Update > Install Update pane.	Delete or install the existing update package.
The wrong update package was automatically downloaded.	If there is more than one new package update available, the most recent one is automatically downloaded.	Delete the update package that was automatically downloaded, and then manually download and install the desired update package.

Manually check for an update package

You can use the PowerProtect Data Manager user interface to manually check for a new update package.

Prerequisites

SupportAssist must be enabled. For more information, see *Configuring SupportAssist for PowerProtect Data Manager* on page 158 and *Enable or disable SupportAssist* on page 161.

Steps

1. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.
2. Click , and then select **Software Update**.
3. In the **Software Update** window, click **Check for Updates**.
4. In the **Software Update > Check for Updates** pane, click **Check Now**.

Results

All available update packages are displayed in the **Software Update > Check for Updates** pane. If there is more than one update package available, more information about each can be seen.

Download an update package

If an update package is available, you can download it to PowerProtect Data Manager for later installation.

About this task

The **Software Update > Install Update** pane displays any update package ready to install. If an update package is already ready to install, it will be overwritten by this procedure.

Steps

1. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.
2. Click , and then select **Software Update**.
3. In the **Software Update** window, click **Check for Updates**.
All available update packages are displayed. If there is more than one update package available, more information about each can be seen.
4. From the **Software Update > Check for Updates** pane, select the update package to download, and then click **Download**.

 **NOTE:** If an update package is already ready to install, you will be asked to confirm if you want to overwrite it.

Results

The update package is downloaded to PowerProtect Data Manager and displayed as ready to install in the **Software Update > Install Update** pane.

Upload an update package

If SupportAssist is not enabled, you do not want to enable it, and you know there is a new update package available that you want to install, then you need to manually upload it to PowerProtect Data Manager.

About this task

The **Software Update > Install Update** pane displays any update package ready to install. If an update package is already ready to install, it will be overwritten by this procedure.

Steps

1. Download the update package from Dell EMC Support Downloads and Drivers.
2. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.
3. Click , and then select **Software Update**.
4. In the **Software Update** window, click **Install Update**.
5. From the **Software Update > Install Update** pane, click **Upload Package**, and select the upload package to upload.

 **NOTE:** If an update package is already ready to install, you will be asked to confirm if you want to overwrite it.

Results

The update package is uploaded to PowerProtect Data Manager and displayed as ready to install in the **Software Update > Install Update** pane.

Delete an update package

If you decide you are not going to install an update package that is ready to install, you can delete it.

Steps

1. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.
2. Click , and then select **Software Update**.
3. In the **Software Update** window, click **Install Update**.
The **Software Update > Install Update** pane displays the update package ready to install.
4. From the **Software Update > Install Update** pane, click **Delete**.

Results

The **Software Update > Install Update** pane is empty, and the update package is no longer ready to install.

Perform a precheck on an update package

After an update package has been downloaded or uploaded, you can perform a precheck on it to verify it is compatible with the current configuration of PowerProtect Data Manager.

Steps

1. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.

2. Click , and then select **Software Update**.
3. In the **Software Update** window, click **Install Update**.
The **Software Update > Install Update** pane displays the update package ready to install.
4. From the **Software Update > Install Update** pane, click **Precheck**.

 **NOTE:** You can navigate away from the **Software Update > Install Update** pane while the precheck is running or even exit the PowerProtect Data Manager interface without interrupting the precheck.

Results

The results of the precheck are displayed in the **Software Update > Install Update** pane.

Install an update package

After an update package has been downloaded or uploaded, you can install it to PowerProtect Data Manager.

Steps

1. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.
2. Click , and then select **Software Update**.
3. In the **Software Update** window, click **Install Update**.
The **Software Update > Install Update** pane displays the update package ready to install.
4. From the **Software Update > Install Update** pane, click **Install**.

Results

If the installation succeeds, the update package is removed from the **Software Update > Install Update** pane. However, you can click the **History** tab from **Software Update > Install Update** pane to see a history of all successful updates.

If the installation fails, the update package remains in the **Software Update > Install Update** pane, along with details of the failure.

Updating the version of PowerProtect Data Manager

This section provides instructions for updating PowerProtect Data Manager from an older version to the most recent version.

If you are updating PowerProtect Data Manager from version 19.8 to version 19.9, follow the steps in [Update PowerProtect Data Manager from version 19.8 to version 19.9 on page 214](#).

If you are updating PowerProtect Data Manager from version 19.7 to version 19.9, follow the steps in [Update PowerProtect Data Manager from version 19.7 to version 19.9 on page 216](#).

If you are updating PowerProtect Data Manager from versions 19.3–19.6 to version 19.9, follow the steps in [Update PowerProtect Data Manager from versions 19.3–19.6 to version 19.9 on page 218](#).

Migrating to SupportAssist

SupportAssist provides automated support capabilities for PowerProtect Data Manager systems. SupportAssist replaces Secure Remote Services (SRS).

If you have configured SRS previously, the PowerProtect Data Manager system automatically migrates SRS to SupportAssist when you update PowerProtect Data Manager.

If you do not have SRS configured, you can configure SupportAssist directly. [Connect to the SupportAssist Enterprise on page 159](#) provides instructions for configuring SupportAssist.

Updating DD or DDVE

If you are updating DD or DDVE software at the same time as PowerProtect Data Manager, ensure that the following sequence of events is followed:

1. Disable all protection policies that use the affected DD or DDVE storage systems.
2. If there are any running jobs that were started by the protection policies, wait for them to complete.
NOTE: Any scheduled replication, cloud tier, or extended replication jobs will continue to run, and might fail during the update.
3. Update the DD or DDVE software.
4. Enable all protection policies that were disabled.

Updating from version 19.8 or earlier in a Kubernetes environment

If you are updating from version 19.8 or earlier and have Kubernetes configured, see the *PowerProtect Data Manager for Kubernetes User Guide* for additional Kubernetes-specific considerations.

CAUTION: If you do not follow a Kubernetes-specific update procedure, PowerProtect Data Manager might fail to function properly.

Update PowerProtect Data Manager from version 19.8 to version 19.9

Use this procedure to update PowerProtect Data Manager from 19.8 to version 19.9 or to apply critical updates.

Prerequisites

- Download the update package from Dell EMC Support Downloads and Drivers.
- Only the Administrator role can perform updates.
- Check for running tasks and cancel them or allow them to complete.
- For on-premise installations, take a manual snapshot of the VM in vCenter, or enable automatic snapshots. Ensure that the vCenter hosting PowerProtect Data Manager is added as an asset source, and that the user account associated with the vCenter host has the following permissions:

Setting	vCenter 6.0 and later required privileges	PowerCLI equivalent required privileges
Global	<ul style="list-style-type: none">o Manage custom attributeso Set custom attributes	<ul style="list-style-type: none">o Global.ManageCustomFieldso Global.SetCustomField
Virtual Machine Snapshot Management	<ul style="list-style-type: none">o Create snapshoto Revert to snapshoto Remove snapshoto Rename snapshot	<ul style="list-style-type: none">o VirtualMachine.State.CreateSnapshoto VirtualMachine.State.RevertToSnapshoto VirtualMachine.State.RemoveSnapshoto VirtualMachine.State.RenameSnapshot

- For cloud-based installations, perform a backup of the AWS instance or Azure VM. The AWS and Azure documentation provides instructions.

About this task

You can update the system by manually downloading update packages or by connecting to a Secure Remote Services (SRS) gateway. When PowerProtect Data Manager is licensed and you have registered the SRS gateway host with PowerProtect Data Manager, you can update using SRS. When an update package is available, the packages are uploaded to the SRS gateway. The appliance checks the SRS gateway once a day for available update packages or you can manually check Dell EMC Support Downloads and Drivers for update packages.

NOTE: If SRS is configured and a critical update is available in the SRS gateway, a notification appears in the UI. You can also download available critical updates that appear in the **Support Site** section of the **Upgrade** page.

An update package can update one or more of the following:

- The PowerProtect Data Manager, including application agent installers stored on the PowerProtect Data Manager virtual machine
- External VM Direct appliance
- Kubernetes support
- PowerProtect Search software
- Remote Cloud Disaster Recovery Server


NOTE: If you have your own SSL certificate that you wish to continue using, the *PowerProtect Data Manager Security Configuration Guide* provides more information.

In PowerProtect Data Manager, the update process automatically stops and removes most running jobs and puts the system into maintenance mode. If server Disaster Recovery is enabled, the system performs a Server DR backup. If automatic snapshots are configured, the update process creates a VM snapshot of the system. If the update fails or is aborted, the system uses the snapshot to roll back to the previous state. Once the system is rolled back or update successfully, the snapshot is automatically deleted.

You can check if the PowerProtect Data Manager system is ready to update by running a manual precheck. Run a manual precheck on page 221 provides more information.

Steps

1. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.

2. Click  and then select **Software Update**.

The **Software Update** window lists any packages that have already been downloaded, in descending date order. If you have registered SRS, the latest available PowerProtect Data Manager update package appears in the **Support Site** section of the window. For any package, you can click the down arrow next to the package name to view details about the contents.

NOTE: When the PowerProtect Data Manager system is configured to automatically check for updates, the download of any newly discovered update package cannot proceed if there is already an update package in local storage.

To download the latest available update package, remove the existing package.

3. If you have registered SRS, in the row for the update package, click **Download**.

If you enabled PowerProtect Data Manager to automatically download update packages in **System Settings > Support > Secure Remote Services**, PowerProtect Data Manager downloads the update package automatically.

When the download is complete, the update package appears in the **Packages** section.

4. If you have not registered SRS and you are using the manual package download method:

- a. Click **Upload Package**.
- b. Browse to the path that contains the update package, select the package, and then click **Open**.
- c. Wait until the package has fully downloaded, and then click **OK**.

5. When the update package status indicates **Available**, click  to start the update wizard.

6. The update wizard consists of 3 stages. Click **Next** at the first two stages, and **Finish** at the last stage.

NOTE: You can also click **Back** or **Cancel** at any stage.

- a. **Precheck.** The update manager runs a precheck.

- If a critical issue is found, the update is cancelled. Fix any issues and run the precheck to ensure that the issue is fixed.
- If non-critical issues are found, Dell EMC recommends that you fix any issues and run the precheck before proceeding with the update.

- b. **Security.** Review the details of the security certificate.

NOTE: You will not see this stage of the wizard if you accepted the security certificate during a previous update.

- c. **Summary.** Review the details of the update.

The update begins. The browser is redirected to the Upgrade Manager UI on port 14443. This action enables you to monitor update progress while the PowerProtect Data Manager components are shutdown for the update.


NOTE: To monitor the update status if the connection to the appliance closes, connect to `https://IP_address_appliance:14443`.

The Upgrade Manager status bar enables you to abort the update, if necessary.

When the update completes successfully, the browser is redirected back to the main PowerProtect Data Manager UI login page.

Results

The **Upgrade** page indicates the status of the update.

- If the update fails, but PowerProtect Data Manager is still running:
 1. Wait for the Upgrade Manager to finish processing.
 2. Click **Return to Dashboard** and log in to view the issue.
 3. Click  and then select **Software Update**.
 4. Expand the package that was installed to view the issue that caused the failure:
 - If one or more core update fail, the status of the update package indicates **Failed**.
 - If all core updates complete, but a VM Direct Engine, the Search Cluster, or another non-core component is still processing, the update package status indicates **Installed (Core)**.
 - If all core updates complete, but a VM Direct Engine, the Search Cluster, or another non-core component fails to update, the update package status indicates **Installed With Errors**.
 5. Fix the issue that caused the failure and run the precheck again.

If the precheck is successful, the package status changes to **Available** and the update can be retried.

6. Retry the update.

When you retry the update, PowerProtect Data Manager only retries the components that failed.

- If the update fails and PowerProtect Data Manager is not running:
 1. Click **Export Logs** to download the log files for troubleshooting.
 2. If an automatic snapshot was taken, click **Rollback to snapshot** to restore the core PowerProtect Data Manager system to its state before the update.
 3. Review the log files to determine the cause of the failure.
 - If you can resolve the issues manually, try the update again.
 - If you cannot resolve the issues, contact Dell EMC Support.

Update PowerProtect Data Manager from version 19.7 to version 19.9

Use this procedure to update PowerProtect Data Manager from 19.7 to version 19.9 or to apply critical updates.

Prerequisites

- Download the update package from Dell EMC Support Downloads and Drivers.
- Only the Administrator role can perform updates.
- Check for running tasks and cancel them or allow them to complete.
- For on-premise installations, take a manual snapshot of the VM in vCenter, or enable automatic snapshots. Ensure that the vCenter hosting PowerProtect Data Manager is added as an asset source, and that the user account associated with the vCenter host has the following permissions:

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
Virtual Machine Snapshot Management	<ul style="list-style-type: none">◦ Create snapshot◦ Revert to snapshot◦ Remove snapshot◦ Rename snapshot	<ul style="list-style-type: none">◦ VirtualMachine.State.CreateSnapshot◦ VirtualMachine.State.RevertToSnapshot◦ VirtualMachine.State.RemoveSnapshot◦ VirtualMachine.State.RenameSnapshot

- For cloud-based installations, perform a backup of the AWS instance or Azure VM. The AWS and Azure documentation provides instructions.

About this task

You can update the system by manually downloading update packages or by connecting to a Secure Remote Services (SRS) gateway. When PowerProtect Data Manager is licensed and you have registered the SRS gateway host with PowerProtect Data Manager, you can update using SRS. When an update package is available, the packages are uploaded to the SRS gateway. The appliance checks the SRS gateway once a day for available update packages or you can manually check Dell EMC Support Downloads and Drivers for update packages.

NOTE: If SRS is configured and a critical update is available in the SRS gateway, a notification appears in the UI. You can also download available critical updates that appear in the **Support Site** section of the **Upgrade** page.

An update package can update one or more of the following:

- The PowerProtect Data Manager, including application agent installers stored on the PowerProtect Data Manager virtual machine
- External VM Direct appliance
- Kubernetes support
- PowerProtect Search software
- Remote Cloud Disaster Recovery Server


NOTE: If you have your own SSL certificate that you wish to continue using, the *PowerProtect Data Manager Security Configuration Guide* provides more information.

In PowerProtect Data Manager, the update process automatically stops and removes most running jobs and puts the system into maintenance mode. If server Disaster Recovery is enabled, the system performs a Server DR backup. If automatic snapshots are configured, the update process creates a VM snapshot of the system. If the update fails or is aborted, the system uses the snapshot to roll back to the previous state. Once the system is rolled back or update successfully, the snapshot is automatically deleted.

You can check if the PowerProtect Data Manager system is ready to update by running a manual precheck. Run a manual precheck on page 221 provides more information.

Steps

1. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.

2. Click  and then select **Software Update**.

The **Software Update** window lists any packages that have already been downloaded, in descending date order. If you have registered SRS, the latest available PowerProtect Data Manager update package appears in the **Support Site** section of the window. For any package, you can click the down arrow next to the package name to view details about the contents.

NOTE: When the PowerProtect Data Manager system is configured to automatically check for updates, the download of any newly discovered update package cannot proceed if there is already an update package in local storage.

To download the latest available update package, remove the existing package.

3. If you have registered SRS, in the row for the update package, click **Download**.

If you enabled PowerProtect Data Manager to automatically download update packages in **System Settings > Support > Secure Remote Services**, PowerProtect Data Manager downloads the update package automatically.

When the download is complete, the update package appears in the **Packages** section.

4. If you have not registered SRS and you are using the manual package download method:

- a. Click **Upload Package**.
- b. Browse to the path that contains the update package, select the package, and then click **Open**.
- c. Wait until the package has fully downloaded, and then click **OK**.

5. When the update package status indicates **Available**, click  to start the update wizard.

6. The update wizard consists of 3 stages. Click **Next** at the first two stages, and **Finish** at the last stage.

NOTE: You can also click **Back** or **Cancel** at any stage.

a. **Precheck.** The update manager runs a precheck.

- If a critical issue is found, the update is cancelled. Fix any issues and run the precheck to ensure that the issue is fixed.


- If non-critical issues are found, Dell EMC recommends that you fix any issues and run the precheck before proceeding with the update.

b. **Authentication.** Enter the lockbox passphrase if required.

Review the section Lockbox passphrase required when updating from some versions on page 221.

c. **Summary.** Review the details of the update.

The update begins. The browser is redirected to the Upgrade Manager UI on port 14443. This action enables you to monitor update progress while the PowerProtect Data Manager components are shutdown for the update.

 **NOTE:** To monitor the update status if the connection to the appliance closes, connect to `https://IP_address_appliance:14443`.


The Upgrade Manager status bar enables you to abort the update, if necessary.

When the update completes successfully, the browser is redirected back to the main PowerProtect Data Manager UI login page.

Results

The **Upgrade** page indicates the status of the update.

- If the update fails, but PowerProtect Data Manager is still running:


1. Wait for the Upgrade Manager to finish processing.
2. Click **Return to Dashboard** and log in to view the issue.
3. Click , and then select **Software Update**.
4. Expand the package that was installed to view the issue that caused the failure:
 - If one or more core update fail, the status of the update package indicates **Failed**.
 - If all core updates complete, but a VM Direct Engine, the Search Cluster, are another non-core component is still processing, the update package status indicates **Installed (Core)**.
 - If all core updates complete, but a VM Direct Engine, the Search Cluster, or another non-core component fails to update, the update package status indicates **Installed With Errors**.
5. Fix the issue that caused the failure and run the precheck again.

If the precheck is successful, the package status changes to **Available** and the update can be retried.

6. Retry the update.

When you retry the update, PowerProtect Data Manager only retries the components that failed.

- If the update fails and PowerProtect Data Manager is not running:

1. Click **Export Logs** to download the log files for troubleshooting.
2. If an automatic snapshot was taken, click **Rollback to snapshot** to restore the core PowerProtect Data Manager system to its state before the update.
3. On the **Upgrade** page, click  to delete the failed update package.
4. Review the log files to determine the cause of the failure.
 - If you can resolve the issues manually, try the update again.
 - If you cannot resolve the issues, contact Dell EMC Support.

Update PowerProtect Data Manager from versions 19.3–19.6 to version 19.9

Use this procedure to update PowerProtect Data Manager from versions 19.3 through 19.6 to version 19.9 or to apply critical updates.

Prerequisites

- Download the update package from Dell EMC Support Downloads and Drivers.
- Only the Administrator role can perform updates.
- Check for running tasks and cancel them or allow them to complete.
- For on-premise installations, take a manual snapshot of the VM in vCenter, or enable automatic snapshots. Ensure that the vCenter hosting PowerProtect Data Manager is added as an asset source, and that the user account associated with the vCenter host has the following permissions:

Setting	vCenter 6.0 and later required privileges	PowerCLI equivalent required privileges
Global	<ul style="list-style-type: none"> o Manage custom attributes o Set custom attributes 	<ul style="list-style-type: none"> o Global.ManageCustomFields o Global.SetCustomField
Virtual Machine Snapshot Management	<ul style="list-style-type: none"> o Create snapshot o Revert to snapshot o Remove snapshot o Rename snapshot 	<ul style="list-style-type: none"> o VirtualMachine.State.CreateSnapshot o VirtualMachine.State.RevertToSnapshot o VirtualMachine.State.RemoveSnapshot o VirtualMachine.State.RenameSnapshot

- For cloud-based installations, perform a backup of the AWS instance or Azure VM. The AWS and Azure documentation provides instructions.
- If you are updating from version 19.3, ensure that you run an ad hoc DR backup operation to back up the Search Service, which is not included in the automatic DR backup that runs before the update.

About this task

You can update the system by manually downloading update packages or by connecting to a Secure Remote Services (SRS) gateway. When PowerProtect Data Manager is licensed and you have registered the SRS gateway host with PowerProtect Data Manager, you can update using SRS. When an update package is available, the packages are uploaded to the SRS gateway. The appliance checks the SRS gateway once a day for available update packages or you can manually check for update packages.

- NOTE:** If SRS is configured and a critical update is available in the SRS gateway, a notification appears in the UI. You can also download available critical updates that appear in the **Support Site** section of the **Software Upgrade** window.

An update package can update one or more of the following:

- The PowerProtect Data Manager software, including application agent installers stored on the PowerProtect Data Manager virtual machine
- External VM Direct appliance
- Kubernetes support
- PowerProtect Search software
- Remote Cloud Disaster Recovery Server


- NOTE:** If you have your own SSL certificate that you wish to continue using, the *PowerProtect Data Manager Security Configuration Guide* provides more information.

In PowerProtect Data Manager, the update process automatically stops most running jobs and puts the system into maintenance mode. If server Disaster Recovery is enabled, the system performs a Server DR backup. If automatic snapshots are configured, the update process creates a virtual machine snapshot of the system. If the update fails or is aborted, the system uses the snapshot to roll back to the previous state. Once the system is rolled back or update successfully, the snapshot is automatically deleted.

In PowerProtect Data Manager 19.5 and later, you can check if the PowerProtect Data Manager system is ready to update by running a manual precheck. Run a manual precheck on page 221 provides more information.

- NOTE:** When you update PowerProtect Data Manager, you are accepting the terms of the latest product EULA. If AutoSupport is enabled, you are also accepting the latest Telemetry Software Terms. It is recommended that you review the Telemetry Software Terms and EULA terms before continuing with the update.
- After you upload the update package, the latest EULA (in both text and PDF format) and the Telemetry Software Terms (in PDF format) are available in the `/data01/brs/update/eulas` folder.

Steps

1. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.
2. Click  and then select **Upgrade**.


The window lists any packages that have already been downloaded, in descending date order. If you have registered SRS, the latest available PowerProtect Data Manager update package appears in the **Support Site** section of the window. For any package, you can click the down arrow next to the package name to view details about the contents.


- NOTE:** When the PowerProtect Data Manager system is configured to automatically check for update, the download of any newly discovered update package cannot proceed if there is already an update package in local storage.

To download the latest available update package, remove the existing package.

3. If you have registered SRS, in the row for the update package, click **Download**.
If you enabled PowerProtect Data Manager to automatically download update packages in **System Settings > Support > Secure Remote Services**, PowerProtect Data Manager downloads the update package automatically.

When the download is complete, the update package appears in the **Packages** section.

4. If you have not registered SRS and you are using the manual package download method:
 - a. Click **Upload Package**.
 - b. Browse to the path that contains the update package, select the package, and then click **Open**.
 - c. Wait until the package has fully downloaded, and then click **OK**.
5. When the update package status indicates **Available**, click  to start the update.


 **NOTE:** In PowerProtect Data Manager versions that are earlier than 19.5, click **Perform Upgrade**.

The **Software Upgrade** wizard appears.

6. On the **Precheck** page, the software update manager runs a precheck:
 - If a critical issue is found, the update is cancelled and the **Next** button is disabled. Fix any issues and then run the precheck again to ensure that the issue is fixed. When the precheck completes successfully, click **Next**.
 - If non-critical issues are found, you can click **Next** to proceed with the update, but Dell EMC recommends that you fix any issues and then run the precheck again before proceeding with the update.
 - If the precheck completes successfully with no issues, click **Next**.
7. On the **Security** page:
 - a. Enter the lockbox passphrase, if required. For example, if updating from PowerProtect Data Manager 19.7 to version 19.9, the lockbox passphrase is required if you have not previously accepted the certificate of the update package.
Review the section **Lockbox passphrase required when updating from some versions** on page 221 to determine if the version you are update from requires you to specify the lockbox passphrase.
 - b. Click **Next**.

8. On the **Summary** page, review the update package and certificate information, and then click **Finish** to proceed.

The update begins. The browser is redirected to the **Upgrade Manager UI** on port 14443. This action enables you to monitor update progress while the PowerProtect Data Manager components are shut down for the update.


 **NOTE:** To monitor the update status if the connection to the appliance closes, connect to `https://IP_address_appliance:14443`.

The Upgrade Manager status bar enables you to abort the update, if necessary.

When the update completes successfully, the browser is redirected back to the main PowerProtect Data Manager UI login page.

Results


The **Software Upgrade** window indicates the status of the update:

- If the update fails, but PowerProtect Data Manager is still running:
 1. Wait for the Upgrade Manager to finish processing.
 2. Click **Return to Dashboard** and log in to view the issue.
 3. Click , and then select **Upgrade**.
 4. Expand the package that was installed to view the issue that caused the failure:
 - o If one or more core update fail, the status of the update package indicates **Failed**.
 - o If all core updates complete, but one or more non-core components, such as vProxies and Search Cluster are still processing, the update package status indicates **Installed (Core)**.
 - o If all core updates complete, but one or more non-core components, such as vProxies and Search Cluster fail to update, the update package status indicates **Installed With Errors**.
 5. Fix the issue that caused the failure and run the precheck again.

If the precheck is successful, the package status changes to **Available** and the update can be retried.

6. Retry the update.

When you retry the update, PowerProtect Data Manager only retries the components that failed.

- If the update fails and PowerProtect Data Manager is not running:
 1. Click **Export Logs** to download the log files for troubleshooting.
 2. If an automatic snapshot was taken, click **Rollback to snapshot** to restore the core PowerProtect Data Manager system to its state before the update.
 3. On the **Upgrade** page, click  to delete the failed update package.
 4. Review the log files to determine the cause of the failure.
 - o If you can resolve the issues manually, try the update again.
 - o If you cannot resolve the issues, contact Dell EMC Support.

Run a manual precheck


For PowerProtect Data Manager versions 19.5 and 19.6, you can run a manual precheck to check if the PowerProtect Data Manager system is ready to update or to verify that any issues that caused a previous precheck to fail are now resolved.

About this task


To run a manual precheck, complete the following steps:

Steps

1. Log in to the PowerProtect Data Manager user interface as a user with the Administrator role.

2. Click  and then select **Upgrade**.

3. To upload an update package:

 **NOTE:** You can skip this step if you have already uploaded the update package.

- a. Click **Upload Package**, browse to the path that contains the update package, select the package, and then click **Open**.
- b. Wait until the package status is Available, and then click **OK**.

Click the down arrow next to the package name to view details about the contents.

4. To run the precheck, click  in the **Actions** column.

When the precheck is complete, a dialog box lists any areas that require attention, such as indication that the update is disruptive or requires a reboot. The dialog box also includes warnings about running tasks or active sessions that should be addressed before the update. Click the links that are provided to go to the management page for the active events, where you can cancel them or allow them to complete before continuing.

The dialog box also indicates if any application agents managed by PowerProtect Data Manager are not compatible with the latest version of the PowerProtect Data Manager system. Manually update the application agents to the latest version before you update the PowerProtect Data Manager system.

If critical issues are found, the precheck fails and the update cannot proceed. If non-critical issues are found, Dell EMC recommends that you fix any issues before proceeding with the update.

Lockbox passphrase required when updating from some versions

In some circumstances, the lockbox passphrase is required to proceed with the PowerProtect Data Manager update.

The following table identifies whether the lockbox passphrase is required when updating to version 19.9, depending on the PowerProtect Data Manager version you are updating from.

Table 52. Lockbox passphrase requirements

Updating from version	Lockbox passphrase required?
19.3	See note.
19.4	No
19.5	No
19.6	No

Table 52. Lockbox passphrase requirements (continued)

Updating from version	Lockbox passphrase required?
19.7	Yes, if you have not previously updated the system and accepted the update-package certificate. Otherwise, no.
19.8	No

- i** **NOTE:** Before automatic lockbox management, updates from version 19.3 required you to provide the lockbox passphrase. Because PowerProtect Data Manager now manages this passphrase, you can type any text in the passphrase field to bypass this step. PowerProtect Data Manager overrides this text with the decrypted lockbox passphrase.

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

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/agentavc` 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 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 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.

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.

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: 10.0.0.1

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/vl/backups` repository. If all 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 clearing 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/vl/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.

Backing Up and Recovering a vCenter Server

Topics:

- Backing up and recovering a vCenter server
- vCenter deployments overview
- Protecting an embedded PSC
- Protecting external deployment models
- vCenter server restore workflow
- Platform Services Controller restore workflow
- Additional considerations
- Command reference

Backing up and recovering a vCenter server

The following sections describe how to protect the vCenter server Appliance (VCSA) and the Platform Services Controllers (PSC). It is intended for virtual administrators who utilize the distributed model of the vCenter server and require protection of the complete vCenter server infrastructure.

vCenter deployments overview

You can protect vCenter 5.5 deployments with PowerProtect Data Manager by using the vProxy appliance. The instructions in this section assume that the vCenter server and the Platform Services Controller (PSC) are deployed as virtual machines.

For the restores to complete successfully:

- Ensure that these virtual machines use a fully qualified domain name (FQDN) with correct DNS resolution.
- Ensure that the host name of the machine is configured as an IP address. Note that if the host name is configured as an IP address, the IP address cannot be changed.

There are mainly two types of vCenter deployments:

- vCenter server Appliance/Windows Virtual Machine with an embedded PSC.
- vCenter server (also multiple) Appliance/Windows virtual machine with an external PSC.

This type has two sub categories:

- vCenter server environment with a single external PSC.
- vCenter server environment with multiple PSC instances. This environment contains multiple vCenter server instances registered with different external PSC instances that replicate their data.

Protecting an embedded PSC

The following section describes backup and recovery options for protecting an embedded PSC.

Backup

You can perform a backup of an embedded PSC by using the following guidelines.

1. Create a protection policy, and then add the vCenter virtual machine to the protection policy.
2. Select the full virtual machine and not individual disks.
3. Run the scheduled or on-demand (ad-hoc) protection policy.

Recovery

Depending on the type of failure, you can perform the virtual machine recovery by using one of the following methods.

- Restore to original — This method is valid only when the vCenter Server Appliance (VCSA) is intact and running, but corrupted.
- Recover as a new virtual machine to a managed ESXi server (Virtual Machine Recovery). Use this method if you have completely lost your VCSA. Note that this vCenter must be registered with PowerProtect Data Manager.
- Direct restore to ESXi server. Direct restore to ESXi will be the main use case.

Direct restore to ESXi

If the virtual machine you protected with PowerProtect Data Manager was a vCenter virtual machine, but this virtual machine and vCenter is now lost or no longer available, direct restore to ESXi enables you to recover the virtual machine directly to an ESXi host without a vCenter server.

Prerequisites

Direct Restore to ESXi restore requires either the embedded VM Direct engine with PowerProtect Data Manager, or an external VM Direct appliance that is added and registered to PowerProtect Data Manager.

Additionally, ensure that you disconnect the ESXi host from the vCenter server.

Steps

1. From the PowerProtect Data Manager UI, select **Restore > Assets**, and then select the **Virtual Machine** tab.
The **Restore** window displays all of the virtual machines available for restore.
 2. Select the checkbox next to the desired virtual machine and click **View Copies**.
 **NOTE:** If you cannot locate the virtual machine, you can also use the filter in the **Name** column to search for the name of the specific virtual machine or click the **File Search** button to search on specific criteria.
- The **Restore > Assets** window provides a map view in the left pane and copy details in the right pane.
- When a virtual machine is selected in the map view, the virtual machine name displays 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 display in the right pane.
3. If the backup is on a DD system, click **DD**, and then select from one of the available copies that display in the table.
 4. In the right pane, select the checkbox next to the virtual machine backup you want to restore, and then click **Direct Restore to ESXi**.
The **Direct Restore to ESXi** wizard appears.
 5. On the **Options** page:
 - a. (Optional) Select **Reconnect the virtual machine's NIC when the recovery completes**, if desired. **Power on the virtual machine when the recovery completes** is selected by default.
 - b. For low-bandwidth environments, select **Enable DDBoost Compression**.
This option reduces network usage by compressing data on the protection storage system before transfer to the VM Direct Engine, which decompresses the data. Compression reduces restore times but increases CPU usage on both systems.
 - c. Click **Next**.
 6. On the **ESX Host Credentials** page:
 - a. In the **ESX Host** field, type the IP of the ESXi server where you want to restore the virtual machine backup.
 - b. Specify the root **Username** and **Password** for the ESXi Server.
 - c. Click **Next**.
 7. On the **Datastore** page, select the datastore where you want to restore the virtual machine disks, and then click **Next**.
 - To restore all of the disks to the same location, keep the **Configure per disk** slider to the left, and then select the datastore from the **Storage** list.
 - To restore disks to different locations, move the **Configure per disk** slider to the right, and then:
 - a. For each available disk that you want to recover, select a datastore from the **Storage** list.
 - b. Select the type of provisioning you want to apply to the disk from the **Disk Format** list.
 8. On the **Summary** page:

- a. Review the information to ensure that the details are correct.
 - b. Click **Restore**.
9. Go to the **Jobs** window to monitor the restore.
A restore job appears with a progress bar and start time.

Protecting external deployment models

Review the backup and recovery options for protecting external deployments.

Backup

You can perform a backup by using the following guidelines:

1. Create one protection policy and add the vCenter virtual machine and PSC virtual machine to the policy. This will ensure that snapshots are taken at the same time.
2. Ensure that you select the full virtual machine and not individual disks.
3. Run the scheduled or on-demand (ad-hoc) protection policy.

NOTE: Ensure that you back up all vCenter server and PSC instances at the same time

Recovery

Depending on the failure, you can perform virtual machine recovery by using one of the following methods:

- Restore to original — This method is valid only when the VCSA is intact and running, but corrupted.
- Recover as a new virtual machine to a managed ESXi server; Use this method if you have completely lost your VCSA. Note that the vCenter where the VCSA resides must be registered with PowerProtect Data Manager.
- Emergency recovery to an ESXi server. For Emergency recovery, perform the steps specified in the section Direct restore to ESXi on page 122.

NOTE: In the event of a complete environment failure, PSC should be restored first, followed by the vCenter server restore.

The following scenarios provide specific instructions based on the number of vCenter server appliances and external PSCs in the environment and the extent of the failure.

vCenter server appliance(s) with one external PSC where PSC fails

Steps

1. Perform an image-level recovery of the PSC by using one of the methods indicated above, and then power ON the virtual machine.
2. Verify that all PSC services are running.
 - For a PSC deployed as an appliance, run the `service-control --status --all` command in the appliance shell.
 - For a PSC installed on Windows, from the Windows Start menu, select **Control Panel > Administrative Tools > Services**.
3. Log into the vCenter server appliance shell as `root`.
4. Verify that no vCenter services are running, or stop any vCenter services that are running by typing `service-control --stop`.
5. Run the `vc-restore` script to restore the vCenter virtual machines.
 - For a vCenter server appliance, type `vc-restore -u psc_administrator_username -p psc_administrator_password`
 - For a vCenter Server installed on Windows, go to `C:\Program Files\VMware\VMware vCenter Server\`, and then run `vc-restore -u psc_administrator_username -p psc_administrator_password` where `psc_administrator_username` is the vCenter Single Sign-On administrator user name, which must be in UPN format.
6. Verify that all vCenter services are running and the vCenter Server is started, as specified in step two.
7. Perform a log in test to the vCenter Server.

If the restore was successful, the login completes successfully.

vCenter server appliance is lost but the PSC remains

Steps

1. Perform an image-level recovery of the lost vCenter server by using one of the following methods, and then power ON.
 - Restore to original — This method is valid only when the VCSA is intact and running, but corrupted.
 - Recover as a new virtual machine to a managed ESXi server — Use this method if you have completely lost your VCSA. Note that this vCenter must be registered with PowerProtect Data Manager.
 - Emergency recovery to an ESXi server.
2. After a successful boot, verify that all services are started.
3. Perform a log in test.

vCenter server appliance with multiple PSCs where one PSC is lost, one remains

Steps

1. Repoint the vCenter instance (insert link) to one of the functional PSC in the same SSO domain.
 - ① **NOTE:** Log in to all vCenter servers one by one to determine which vCenter log in fails. This will be the vCenter that requires the repoint steps.
2. Run the following command on the vCenter server appliance:
`cmsso-util repoint --repoint-psc psc_fqdn_or_static_ip [--dc-port port_number]`
 - ① **NOTE:** The square brackets enclose the command options.
3. Perform a log in test on the vCenter server.
4. Deploy the new PSC and join to an active node in the same SSO and site, replacing lost ones.
5. Repoint the vCenter server to the new PSC.

vCenter server appliance remains but all PSCs fail

About this task

① **NOTE:** In this scenario, none of the vCenter logins (SSO user) have been successful.

Steps

1. Restore the most recent PSC backup and wait for the vCenter services to start.
2. Log in to the vCenter server appliance's shell as `root`.
3. Verify that no vCenter services are running, or stop vCenter services.
4. Run the `vc-restore` script to restore the VCSA (refer above for detailed steps).
 - ① **NOTE:** If the login test to any vCenter server appliance fails, then the restored PSC is not the PSC that the vCenter server appliance is pointing to, in which case you may be required to perform a repoint, as described above.
5. Deploy the new PSC and join to an active node in the same SSO domain and site.
6. Repoint vCenter connections as required.

vCenter server appliance remains but multiple PSCs fail

Steps

1. Restore one PSC.

2. Test the vCenter server appliance login. If the login fails, repoint the vCenter server appliance to an active PSC.
3. Deploy the new PSC and join to an active node in the same SSO domain and site.

vCenter server appliance fails

About this task

NOTE: If a total failure has occurred (all PSCs and all vCenter server appliances failed), restore one PSC first before restoring the vCenter server appliance.

Steps

1. Perform an image-level restore of the lost vCenter server by using one of the following methods, and then power ON the vCenter.
 - Restore to original — This method is valid only when the vCenter server appliance is intact and running, but corrupted.
 - Recover as a new virtual machine to a managed ESXi server — Use this method if you have completely lost your vCenter server appliance. Note that this vCenter must be registered with PowerProtect Data Manager.
 - Emergency recovery to an ESXi server.
2. After a successful boot, verify that all vCenter services have started.
3. Perform a log in test.
4. If the log in test fails, then this vCenter server appliance is pointing to an inactive PSC. Repoint to an active node.

vCenter server restore workflow

The following diagram shows the restore workflow for a vCenter server.

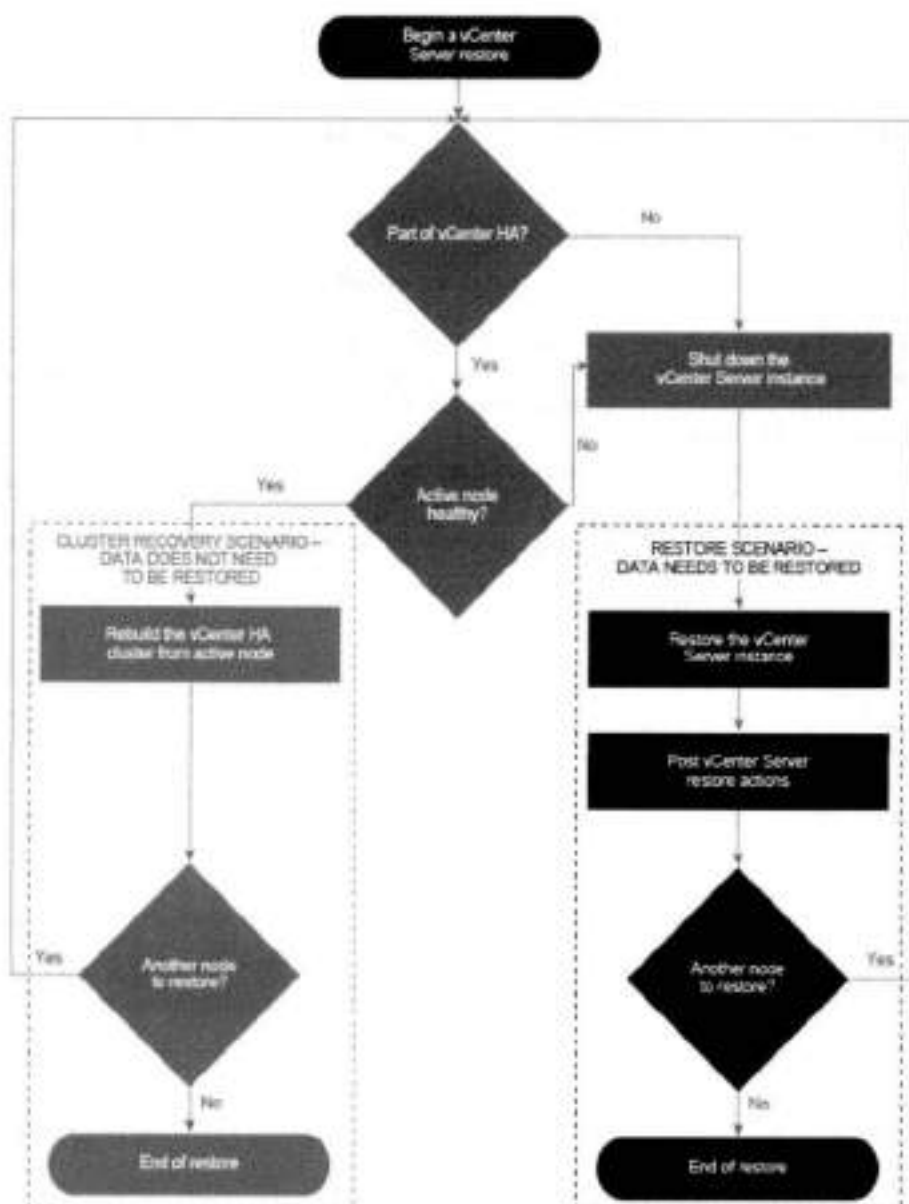


Figure 14. vCenter server restore workflow

Platform Services Controller restore workflow

The following diagram shows the restore workflow for a Platform Services Controller (PSC).

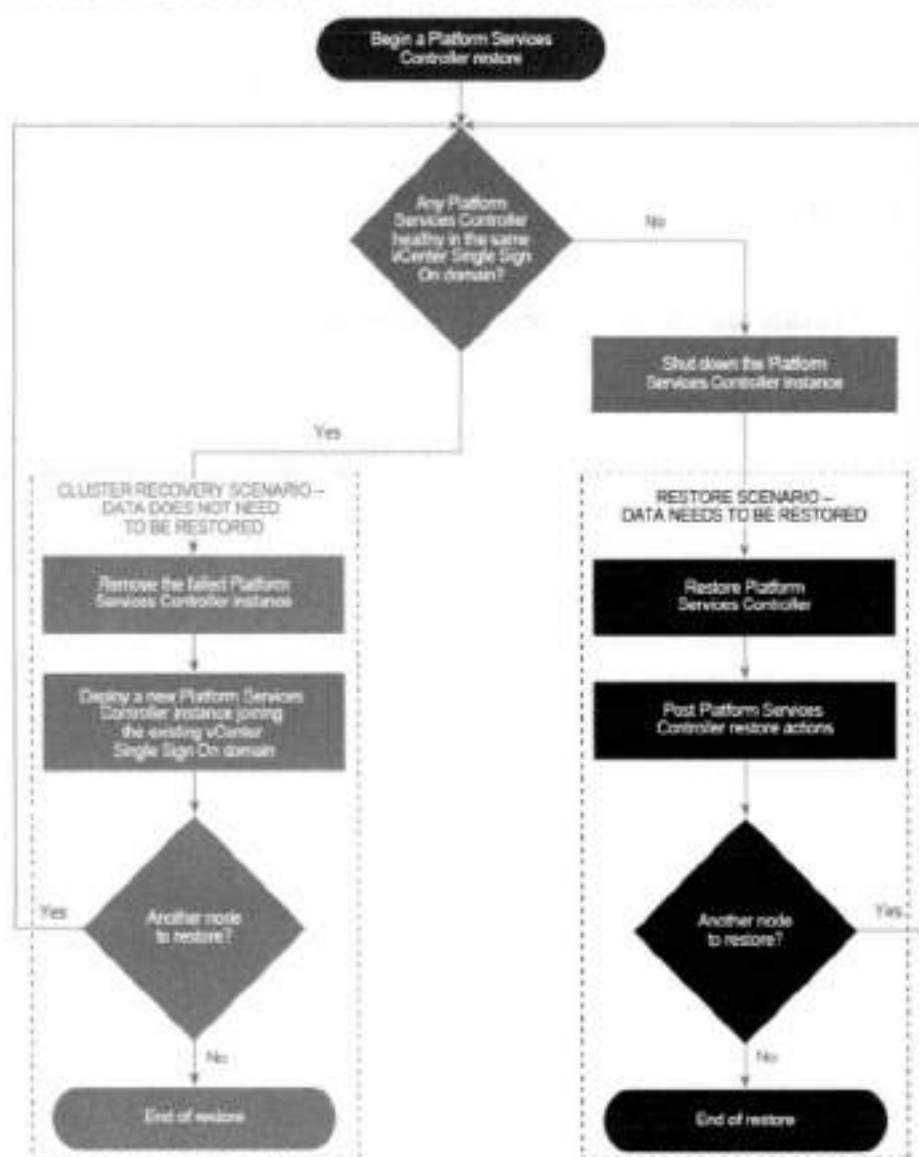


Figure 15. PSC restore workflow

Additional considerations

Review the following additional considerations when backing up and restoring the vCenter server and PSC.

- Backing up the vCenter server will not save the Distributed switch (vDS) configuration as it is stored on the hosts. As a best practice, back up the vDS configuration by using a script that can be used after restoring the virtual center.
- After restoring the PSC, verify that replication has been performed as designed by using the following commands to display the current replication status of a PSC and any of the replication partners of the PSC:
 - For VCSA, go to `/usr/lib/vmware-vmtoolsd/bin` and type `./vdsrepadmin -f showpartnerstatus -h localhost -u administrator -w Administrator_Password`
 - For Windows, open a command prompt and type `cd "%VMWARE_CIS_HOME%\vmtoolsd\`
- For the vCenter server or PSC, do not select advanced quiesce-based backup options. Selecting these options will result in application quiescing on virtual machines, which impacts the overall environment due to stuning.

The VMware vCenter server documentation, available at <https://docs.vmware.com/en/VMware-vSphere/index.html>, provides more information about the vCenter server and PSC.

Command reference

Use the following command to start or stop services in the vCenter server/PSC, or obtain the status:

```
service-control -status/start/stop -all
```

You can use other Replication topology commands, as in the following example.

Replication topology command

```
/usr/lib/vmware-vmdir/bin/vdorepadmin -f showpartners -h localhost -u PSC_Administrator -w password
```

① **NOTE:** You can replace **localhost** with another PSC FQDN to obtain all of the partnerships in the current vSphere domain.

Backing Up VMware Cloud Foundation (VCF) on VxRail

Topics:

- Backing up VCF on VxRail
- VCF and VxRail overview
- VCF components and backup methods
- Check VMware certification
- Backup prerequisites
- The backup script
- Quick protection
- Selective protection: SDDC and NSX-T Managers
- Selective protection: vCenter servers
- Selective protection: vRSLCM, VxRail Manager, Workspace ONE Access, and vRealize Suite virtual machines
- SFTP password change: SDDC and NSX-T Managers
- SFTP password change: vCenter servers
- Backup-script troubleshooting

Backing up VCF on VxRail

The following sections describe how to protect VMware Cloud Foundation (VCF) on VxRail by using a PowerProtect Data Manager command-line backup script.

NOTE: VxRail is the preferred Dell EMC platform for VCF. However, environments that use other VMware-supported vSAN Ready Nodes are also supported by Dell EMC. The following sections also apply to those environments.

VCF and VxRail overview

VCF integrates a VMware cloud infrastructure with cloud management services by using the vRealize software suite to run enterprise applications. The VCF infrastructure is managed by the SDDC Manager, and it includes vSphere compute, vSAN storage, NSX networking, and a range of security implementations.

Dell EMC VxRail is an all-in-one solution that uses Dell EMC PowerEdge servers and its own VxRail hyperconverged infrastructure (HCI) software to provide a fully functional VCF environment to enterprise customers.

For more information about VCF and VxRail, see the following resources:

- The VMware Cloud Foundation documentation
- The *Dell EMC VxRail Administration Guide* at Customer Support
- About VMware Cloud Foundation on Dell EMC VxRail

VCF components and backup methods

Understanding the backup method used by a VCF component aids in understanding how the VCF component is protected by the backup script. The following tables show the VCF components of the different backup methods.

Table 53. VCF components of file-based backups

Backup Method	Component
File based	NSX-T Data Center
	SDDC Manager
	vCenter Server

- Assets of these components are first copied to an external server that uses secure file transfer protocol (SFTP) or another supported protocol. After that, the external server is backed up by PowerProtect Data Manager.
- If using quick protection, these components are automatically protected.

Table 54. VCF components of image-based backups

Backup Method	Component	Automatically discovered
Image based	vRealize Suite Lifecycle Manager (vRSLCM)	VCF 4.0
	vRealize Automation	VCF 4.1
	vRealize Business	No
	vRealize Log Insight	VCF 4.1
	vRealize Network Insight	No
	vRealize Operations Manager	VCF 4.1
	VxRail Manager	No
	Workspace ONE Access	VCF 4.1

- Assets of these components are backed up directly by PowerProtect Data Manager.
- The **Automatically discovered** column displays the minimum required version of VCF for a component to be automatically discovered, as well as those components that are not automatically discovered by any version of VCF.
- If using quick protection, the automatically discovered components are automatically protected.

All image-based backups follow the VMware quiescing recommendations for VCF virtual machines that are part of VMware Validated Design (VVD):

Table 55. VCF components and quiescing

Component	Quiescing
vRealize Suite Lifecycle Manager	Enabled
Workspace ONE Access	Enabled
vRealize Log Insight	Disabled
vRealize Operations Manager	Disabled
vRealize Automation	Enabled

Check VMware certification

Use this method to check the versions of PowerProtect Data Manager that VMware has certified to work with their products.

About this task

VMware certification allows customers to receive support from VMware for any VMware-specific features related to PowerProtect Data Manager.

NOTE: VMware will only certify a version of PowerProtect Data Manager after it has been released and tested. If you are waiting for the current version of PowerProtect Data Manager to be certified, you can continue to check its status.

Steps

1. In a browser, navigate to the VMware Compatibility Guide.
2. Select **All > Dell EMC > All**.
3. Click **Update and View Results**.
4. In the **Solution Name** column, look for *EMC PowerProtect Data Manager* entries.
5. Review the information in the corresponding **Solution Version** and **Supported Releases** columns.

Backup prerequisites

Ensure the following prerequisites are met before backing up VCF on VxRail:

- VCF is at a supported version. For more information, see the PowerProtect Data Manager compatibility matrix at the E-Lab Navigator.
- Any external server (using SFTP or another supported protocol) used in a file-based backup has been discovered as a File System asset in PowerProtect Data Manager.
- Any vCenter server being protected has been added as an asset source in PowerProtect Data Manager.
- PowerProtect Data Manager and the vCenter, SDDC, and NSX-T managers are all set to the same time zone and have their clocks synchronized.
- PowerProtect Data Manager and VCF do not have backup schedules that would back up the same assets at the same time.
- A VM Direct Engine exists.
- Any backup directory path specified by an external server in a file-based backup exists.
- All credentials provided during the execution of the backup script resolve to accounts with the required permissions to access the related resources.

The backup script

You use a PowerProtect Data Manager script to protect VCF components.

The script is accessible from the PowerProtect Data Manager command line. It provides a series of guided procedures that automate multiple backup operations into a single process. The script can also be used to change external SFTP passwords.

NOTE: This script only backs up the data of protected VCF components. It cannot be used to restore any of the data that is backed up. To restore the data, use the PowerProtect Data Manager and VMware user-interface tools. Ensure that you restore VCF-management data to components in a manner supported by VMware. For more information, go to the VMware Validated Design Documentation website and review the backup and restore procedures of the documentation that corresponds to your version of VCF.

Quick protection

This procedure uses default backup settings and values to protect all VCF components at once. Every vCenter server and any automatically discovered VCF component will be protected. Quick protection requires the least amount of input, but also provides the least amount of choice. For information about the default settings and values used, review the selective-protection procedures that follow.

Steps

1. From a PowerProtect Data Manager command line, type the following two commands:

```
cd /usr/local/brs/lib/synggr/bin
./ppdm-vcf-component-protection.sh
```

2. Enter credentials for the PowerProtect Data Manager server.
3. Enter the IP address or fully qualified domain name (FQDN) of the SDDC Manager server, and then enter credentials for the SDDC Manager server.
4. From the backup-script main menu, enter **1**.

NOTE: Quick protection uses the same external SFTP server and backup schedule for both the SDDC Manager and vCenter servers. It also overrides the existing backup configurations of the SDDC and vCenter servers without prompting.

5. Enter the address of an external SFTP server, including the backup directory path, followed by credentials to access the server. The external SFTP server is also used for vCenter server configuration.

The external SFTP server and backup directory path uses the format **sftp://server_ip_address:port_number/folder/subfolder**. For example:

```
sftp://172.17.62.201:22/upload/backup
```

6. Enter the encryption passphrase for SDDC Manager backups.

The encryption passphrase must be between 12 and 32 characters in length and contain at least two lowercase letters, two uppercase letters, two numbers, and a special character.

NOTE: The encryption passphrase is also used for vCenter server backups, and is required when restoring data. Store the passphrase in a secure location that is separate from the backup files and VCF environment you are protecting.

7. Confirm if common credentials should be used.
 - Enter **y** to provide common credentials for all vCenter servers.
 - Enter **n** to be prompted for the credentials for each individual server.

8. Select the days of the week a backup takes place, and then enter the time of day.

Type a number that represents a day of the week, where **1** represents Sunday. If selecting multiple days of the week, separate the numbers with a space. For example, to select Sunday and Monday:

```
1 2
```

The time of day uses the format **HH:MM** in 24-hour notation. For example, to enter 1:25 p.m.:

```
13:25
```

9. Select both a File System and Virtual Machine protection policy to use.
 - If a default protection policy of either type does not exist, it will be automatically created with a frequency of DAILY, a time of 8:00 PM to 6:00 AM, and a retention of 7 days.
 - A protection policy with the name *VCF-Image-Based-Protection* is used as the default image-based protection policy.
 - A protection policy with the name *VCF-File-Based-(SFTP)-Protection* is used as the default file-based protection policy.
 - If a default protection policy has just been automatically created and it is the only protection policy of that type, it will be automatically used.
 - If a default protection policy already exists, confirm if it should be used or if the protection policy to use should be selected from a list.
10. Enter the IP address or FQDN of any image-based VCF component that is not automatically discovered and that you want to protect. For a list of components that are not automatically discovered, see *VCF components and backup methods* on page 236.

Results

You can monitor the progress of the backup script as it protects the VCF components.

Selective protection: SDDC and NSX-T Managers

This procedure protects just the SDDC and NSX-T manager file-based VCF components, while providing more control over the backup settings used for them than quick protection. To protect other VCF components, refer to the other selective-protection procedures.

Steps

1. From a PowerProtect Data Manager command line, type the following two commands:

```
cd /usr/local/brs/lib/sysmgr/bin
./ppdm-vcf-component-protection.sh
```

2. Enter credentials for the PowerProtect Data Manager server.
3. Enter the IP address or fully qualified domain name (FQDN) of the SDDC Manager server, and then enter credentials for the SDDC Manager server.
4. From the backup-script main menu, enter **2**, and then **1**.
5. To override an existing SDDC Manager backup configuration, enter **y**.
6. To add or modify SDDC Manager backup configuration information, enter the address of an external SFTP server, including the backup directory path, followed by credentials to access the server.

The external SFTP server and backup directory path uses the format `sftp://server_ip_address:port_number/folder/subfolder`. For example:

```
sftp://172.17.62.201:22/upload/backup
```

7. Enter the encryption passphrase for SDDC Manager backups.

The encryption passphrase must be between 12 and 32 characters in length and contain at least two lowercase letters, two uppercase letters, two numbers, and a special character.

NOTE: The encryption passphrase is required when restoring data. Store this passphrase in a secure location that is separate from the backup files and VCF environment you are protecting.

8. The default SSH fingerprint of the external SFTP server is displayed. Confirm that it should be used, or enter a new one.

NOTE: With quick protection, the default SSH fingerprint of the external SFTP server is always used.

9. Select the backup frequency. If you select **HOURLY**, enter the minute of each hour a backup takes place. If you select **WEEKLY**, select the days of the week a backup takes place, and then enter the time of day.

For a weekly backup frequency, type a number that represents a day of the week, where **1** represents Sunday. If selecting multiple days of the week, separate the numbers with a space. For example, to select Sunday and Monday:

```
1 2
```

The time of day uses the format **HH:MM** in 24-hour notation. For example, to enter 1:25 p.m.:

```
13:25
```

10. Enter the backup-retention values described in the following table. The values automatically used by quick protection are also listed.

Table 56. Backup-retention values

Parameter	Value range	Quick-protection default value
Days of daily backups to retain	0–30	7
Days of hourly backups to retain	0–14	7
Backup files to retain	1–600	15
Take backups on state change	Yes or no	Yes

11. Confirm if a new File System protection policy should be created in order to protect the external SFTP server.

- Enter **y** to provide details of the new protection policy.
- Enter **n** to either select from a list of existing protection policies or skip protection of the external SFTP server.

Results

You can monitor the progress of the backup script as it protects the selected VCF components.

Selective protection: vCenter servers

This procedure protects just the vCenter server file-based VCF components, while providing more control over the backup settings used for them than quick protection. To protect other VCF components, refer to the other selective-protection procedures.

Steps

1. From a PowerProtect Data Manager command line, type the following two commands:

```
cd /usr/local/brs/lib/syasmgr/bin
./ppdm-vcf-component-protection.sh
```

2. Enter credentials for the PowerProtect Data Manager server.
3. Enter the IP address or fully qualified domain name (FQDN) of the SDDC Manager server, and then enter credentials for the SDDC Manager server.
4. From the backup-script main menu, enter **2** twice.
5. Select the automatically discovered vCenter servers to protect.

Enter **a** to protect all the servers. Otherwise, enter the numbers that correspond to the individual servers to protect, separating each number with a space.

6. Enter the address of an external server, including the backup directory path, followed by credentials to access the server. Supported protocols for the external server are FTP, SFTP, FTPS, HTTP, HTTPS, NFS, and SMB. The external server and backup directory path uses the format `protocol://server_ip_address:port_number/folder/subfolder`. For example:

```
sftp://172.17.62.201:22/upload/backup
```

7. Select the days of the week a backup takes place, and then enter the time of day.

Type a number that represents a day of the week, where **1** represents Sunday. If selecting multiple days of the week, separate the numbers with a space. For example, to select Sunday and Monday:

```
1 2
```

The time of day uses the format **HH:MM** in 24-hour notation. For example, to enter 1:25 p.m.:

```
13:25
```

8. Confirm if the backups should be encrypted. If they should be encrypted, enter an encryption password. If you enter an encryption password, it must be between 8 and 20 characters in length and contain at least one lowercase letter, one uppercase letter, one number, and one special character.
9. Confirm if historical data should be backed up and the number of backups to retain.

i **NOTE:** In quick protection, the default is to back up historical data and retain all backups.

10. Confirm if common credentials should be used.

- Enter **y** to provide common credentials for all vCenter servers.
- Enter **n** to be prompted for the credentials for each individual server.

11. If there is an existing vCenter server backup configuration, confirm if it should be overridden.

i **NOTE:** Should the existing backup configuration fail to be overridden, the vCenter server will be left without a backup configuration.

12. Confirm if a new File System protection policy should be created in order to protect the external server.

- Enter **y** to provide details of the new protection policy.
- Enter **n** to either select from a list of existing protection policies or skip protection of the external server.

Results

You can monitor the progress of the backup script as it protects the selected VCF components.

Selective protection: vRSLCM, VxRail Manager, Workspace ONE Access, and vRealize Suite virtual machines

This procedure protects all of the image-based VCF components, while providing more control over the backup settings used for them than quick protection. The components protected include vRSLCM, VxRail Manager, Workspace ONE Access, and vRealize Suite virtual machines. To protect file-based VCF components, refer to the other selective-protection procedures.

Steps

1. From a PowerProtect Data Manager command line, type the following two commands:

```
cd /usr/local/brs/lib/symgr/bin
./ppdm-vcf-component-protection.sh
```

2. Enter credentials for the PowerProtect Data Manager server.
3. Enter the IP address or fully qualified domain name (FQDN) of the SDDC Manager server, and then enter credentials for the SDDC Manager server.
4. From the backup-script main menu, enter **2**, and then **3**.
5. Select an image-based VCF component type to protect.

NOTE: You can only select a single component type. To protect more than one component, follow the selective protection steps for each component.

- If you select vRSLCM, select a discovered vRSLCM server to protect.
 - If you select any other component type, enter the IP address or fully qualified domain name (FQDN) of the server to protect.
6. Confirm if a new Virtual Machine protection policy should be created in order to protect the component.
 - Enter **y** to provide details of the new protection policy.
 - Enter **n** to select from a list of existing protection policies.

Results

You can monitor the progress of the backup script as it protects the selected VCF component.

SFTP password change: SDDC and NSX-T Managers

While using the backup script to protect VCF components, you might want to change the password of the external SFTP server account associated with the SDDC and NSX-T Managers.

Steps

1. From a PowerProtect Data Manager command line, type the following two commands:

```
cd /usr/local/brs/lib/symgr/bin
./ppdm-vcf-component-protection.sh
```

2. Enter credentials for the PowerProtect Data Manager server.

3. Enter the IP address or fully qualified domain name (FQDN) of the SDDC Manager server, and then enter credentials for the SDDC Manager server.
4. From the backup-script main menu, enter **3**, and then **1**.
5. Confirm if you want to change the password of the external SFTP server account.
 - Enter **y** to change the password, and then perform the following actions:
 - a. Enter the new password.
 - b. Enter **y** to confirm if the automatically generated SSH fingerprint should be used. Otherwise, enter **n** to provide your own SSH fingerprint.
 - Enter **n** to skip the password change.

Results

You can monitor the progress of the backup script as it changes the password of the external SFTP server account associated with the SDDC and NSX-T managers.

SFTP password change: vCenter servers

While using the backup script to protect VCF components, you might want to change the password of an external SFTP server associated with an automatically discovered vCenter server.

Steps

1. From a PowerProtect Data Manager command line, type the following two commands:


```
cd /usr/local/brs/lib/sysexec/bin
./ppdm-vcf-component-protection.sh
```
2. Enter credentials for the PowerProtect Data Manager server.
3. Enter the IP address or fully qualified domain name (FQDN) of the SDDC Manager server, and then enter credentials for the SDDC Manager server.
4. From the backup-script main menu, enter **3**, and then **2**.
5. Confirm if common credentials should be used.
 - Enter **y** to provide common credentials for all vCenter servers.
 - Enter **n** to be prompted for the credentials for each individual server.
6. Confirm if you want to provide a backup encryption password. This password will be used when backing up the VCF components of all vCenter servers.
7. For each automatically discovered vCenter server, confirm if you want to change the password of the external SFTP server account associated with it.

Results

You can monitor the progress of the backup script as it changes the passwords of all external SFTP server accounts associated with the selected vCenter servers.

Backup-script troubleshooting

The following table provides common error codes and messages, along with explanations or recommended areas of investigation to resolve the problem.

Table 57. Error codes and explanations

Error code or message	Explanation or area of investigation
INVALID_ENCRYPTION_PASSPHRASE Provided encryption passphrase <passphrase> is invalid.	The encryption passphrase specified for external SFTP server is invalid.
Validate Backup Location Details FAILED	The backup location specified for the external SFTP server in the SDDC Manager backup configuration does not exist.
INPUT_PARAM_ERROR Failed to establish SFTP connection to <SFTP server> with username <username> on port <port>.	The credentials specified for the external SFTP server in the SDDC Manager backup configuration are incorrect.
INVALID_ARGUMENT The entered backup password does not adhere to the password requirements.	The encryption passphrase specified in the vCenter server backup configuration is invalid.
INVALID_ARGUMENT Plugin error occurred. Access to the backup server is denied. Check your credentials.	The password specified for the external server in the vCenter server backup configuration is incorrect.
UNAUTHENTICATED Authentication required. com.vmware.vapi.endpoint.method.authentication.required	The credentials specified for the vCenter server are incorrect.
Perform validations for backup server fingerprint FAILED	The SSH fingerprint specified for the external SFTP server in the SDDC Manager backup configuration is invalid.
SCHEDULING_SDDC_MANAGER_BACKUPS_FAILED_REASON_UNKNOWN Unexpected error occurred. Provided backup schedule not applied.	Check for errors on the SDDC Manager.

Table 57. Error codes and explanations (continued)

Error code or message	Explanation or area of investigation
<p>LOCK_NOT_AVAILABLE</p> <p>Lock is not available - SDDC Manager DEPLOYMENT lock to perform Backup & Restore operation.</p>	<p>There are too many pending SDDC Manager jobs. Try running the backup script at another time.</p>
<p>503</p> <p>The data store service is not available. Try again later.</p> <p>remediation timestamp <timestamp> path /api/v2/assets</p>	<p>PowerProtect Data Manager assets cannot currently be queried. Try running the backup script at another time.</p>
<p>503</p> <p>The service is not available. Try again later.</p> <p>remediation timestamp <timestamp> path /api/v2/protection-policies</p>	<p>Protection policies cannot currently be queried. Try running the backup script at another time.</p>

Best Practices and Troubleshooting

Topics:

- Base 10 standard used for size calculations in the PowerProtect Data Manager UI
- Best practices and additional considerations for the VM Direct Engine
- Best practices for vCenter Server backup and restore
- Changing the vCenter server FQDN
- Monitoring storage capacity thresholds
- Replacing security certificates
- Restarting PowerProtect Data Manager
- Scalability limits for vCenter Server, VM Direct Engine and DD systems
- Troubleshooting network setup issues
- Troubleshooting PowerProtect agent service installations
- Troubleshooting PowerProtect agent service operations
- Troubleshooting PowerProtect Data Manager software updates
- Troubleshooting storage units
- Troubleshooting virtual machine backup issues
- Troubleshooting virtual machine restore issues
- Troubleshooting vSphere Plugin deployments

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

For size calculations (for example, asset size, the available space on storage systems), 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.

Best practices and additional considerations for the VM Direct Engine

Review the following information for recommendations and best practices when adding a VM Direct protection engine in PowerProtect Data Manager.

Change the limit of instant access sessions

For DDOS versions 6.2 and higher, PowerProtect Data Manager uses the limit that the DD storage appliance reports, and manages concurrent instant access sessions based on the reported limit.

You can change the limit by modifying a configuration file to override the default value. Note that sessions that exceed the maximum concurrent sessions that are supported are canceled and retried. To change the number of concurrent sessions manually to match the capability of the underlying storage appliance, perform the following steps.

1. Log in to the PowerProtect Data Manager UI as a user with the Administrator role.
2. If not already created, create an `application.yml` file in the `/usr/local/brs/lib/vmdm/config/` directory.

NOTE: The structure of this file requires that you separate fields into individual categories and sub categories, as shown in the following step.

3. In the `application.yml` file, change the instant access session parameter value to override the default value. For example:

```
recovery:
  queue:
    ia_session_allowance: 32
```

4. Run `vmdm stop` and then `vmdm start` to restart the `vmdm` service.

NOTE: Ensure that no other virtual machine operations are running, such as protection and recovery.

Configuring a backup to support vSAN datastores

Backup and recovery functionality is supported for vSAN virtual machines.

When performing backups or restores of virtual machines residing on vSAN datastores, it is highly recommended to deploy the VM Direct appliance on a vSAN datastore. A VM Direct appliance deployed on any one vSAN datastore can be used for backing up virtual machines from other vSAN or non-vSAN datastores by using **Hot Add** or **nbdssl** transport modes, as applicable.

Configuration checklist for common issues

The following configuration checklist provides best practices and troubleshooting tips that might help resolve some common issues.

Basic configuration

Review the following basic configuration requirements:

- Synchronize system time between vCenter and ESX/ESXi/vSphere.
- Assign IPs carefully — do not reuse any IP addresses.
- Use Fully Qualified Domain Names (FQDNs) where possible.
- For any network related issue, confirm that forward and reverse DNS lookups work for each host in the datazone.

Virtual machine configuration

Review the following virtual machine configuration requirements:

- Ensure that the virtual machine has access to and name resolution for the protection storage.
- Ensure that the virtual machine firewall has port rules for the protection storage.
- For application-aware backups, ensure that Microsoft SQL Server instances are enabled for data protection using a SYSTEM account, as described in the section "Microsoft application agent for SQL Server application-aware protection" of the *PowerProtect Data Manager for Microsoft Application Agent SQL Server User Guide*.

Disable vCenter SSL certificate validation

If the vCenter's SSL certificate cannot be trusted automatically, a dialog box appears when adding the vCenter Server as an asset source in the PowerProtect Data Manager UI, requesting certificate approval. It is highly recommended that you do not disable certificate enforcement.

If disabling of the SSL certificate is required, you can perform the following procedure.

CAUTION: These steps should only be performed if you are very familiar with certificate handling and the issues that can arise from disabling a certificate.

1. Create a file named `cbs_vmware_connection.properties` in the `/home/admin` directory on the PowerProtect Data Manager appliance, with the following contents:

```
cbs.vmware_connection.ignore_vcenter_certificate=true
```

2. If not already created, create an `application.yml` file in the `/usr/local/brs/lib/vmdm/config/` directory.

NOTE: The structure of this file requires that you separate fields into individual categories and sub categories, as shown in the following step.

3. In the `application.yml` file, add the following contents:

```
vmware_connection:
  ignore_vcenter_cert: true

discovery:
  ignore_vcenter_cert: true
```

4. Run `cbs stop` to stop the `cbs` service, and then `cbs start` to restart the service.
5. Run `vmdm stop` to stop the `vmdm` service, and then `vmdm start` to restart the service.
6. Perform a test to determine if SSL certificate disabling was successful by adding a vCenter Server using the vCenter's IP address (if the SSL certificate uses FQDN), and then verify that the asset source was added and virtual machine discovery was successful.

File-level restore and SQL restore limitations

This section provides a list of limitations that apply to file-level restore and individual SQL database end instance restore.

Consider the following:

- The VM Direct **FLR Agent** is installed automatically on the target virtual machine for file-level restore when a disk mount operation is initiated. However, if the user does not have sufficient administrator privileges, the mount fails and the **FLR Agent** is not installed. Ensure that the user performing file-level restore is a system administrator. Note that adding a user to the Administrators group does not grant this user sufficient privileges to perform this operation.
- When performing a file-level restore, VMDKs fail to mount with the following error if the **FLR Agent** service is not running on the target virtual machine: "Cannot connect to vProxy Agent: dial tcp <127.0.0.1:<port>: connectex: No connection could be made because the target machine actively refused it."
- If you no longer require the VM Direct **FLR Agent** on the target virtual machine, the agent must be properly uninstalled. If you manually delete VM Direct FLR Agent files instead of uninstalling the agent, and at some point reinstall the agent, subsequent mount attempts to perform restores will fail.

To uninstall the VM Direct **FLR Agent** on Linux:

1. Execute the following command: `/opt/emc/vproxys/bin/preremove.sh`.
2. Uninstall FLR agent package by running `rpm -e emc-vProxy-FLRAgent`.
3. If the uninstall fails due to a broken installation or other issue, you can force removal of the package by running `rpm -e --force emc-vProxy-FLRAgent`.

To uninstall the VM Direct **FLR Agent** on Windows:

1. Select **Control Panel > Programs > Programs and Features**.
 2. Locate **EMC VM Direct FLR**.
 3. Right-click the program and select **Uninstall**.
- When a file-level restore or SQL restore operation is in progress on a virtual machine, no other backup or recovery operation can be performed on this virtual machine. Wait until the file-level restore session completes before starting any other operation on the virtual machine.
 - Clean up from a suspended or canceled mount operation requires a restart of the virtual machine before you can initiate a new mount for the file-level restore.
 - When you enable Admin Approval Mode (AAM) on the operating system for a virtual machine (for example, by setting `Registry/FilterAdministratorToken` to **1**), the administrator user cannot perform a file-level restore to the end user's profile, and an error displays indicating "Unable to browse destination." For any user account control (UAC) interactions, the administrator must wait for the mount operation to complete, and then access the backup folders located at `C:\Program Files (x86)\EMC\vProxy FLR Agent\flr\mountpoints` by logging into the guest virtual machine using Windows Explorer or a command prompt.
 - When you perform file-level restore on Windows 2012 R2 virtual machines, the volumes listed under the virtual machine display as "unknown." File-restore operations are not impacted by this issue.
 - When you perform file-level restore on Ubuntu/Debian platforms, you must enable the root account in the operating system. By default, the root account will be in locked state.
 - You can only restore files and/or folders from a Windows backup to a Windows machine, or from a Linux backup to a Linux machine.

- You must install VMware Tools version 10 or later. For best results, ensure that all virtual machines run the latest available version of VMware Tools. Older versions are known to cause failures when you perform browse actions during file-level restore or SQL restore operations.
- You can perform file-level restore across vCenters as long as the vCenters are configured in PowerProtect Data Manager, and the source and target virtual machine have the same guest operating system. For example, Linux to Linux, or Windows to Windows.
- File-level restore does not support the following virtual disk configurations:
 - LVM thin provisioning
 - Unformatted disks
 - FAT16 file systems
 - FAT32 file systems
 - Extended partitions (Types: 05h, 0Fh, 85h, C5h, D5h)
 - Two or more virtual disks mapped to single partition
 - Encrypted partitions
 - Compressed partitions
- File-level restore of virtual machines with Windows dynamic disks is supported with the following limitations:
 - The restore can only be performed when recovering to a virtual machine different from the original. Also, this virtual machine cannot be a clone of the original.
 - The restore can only be performed by virtual machine administrator users.
 - If Windows virtual machines were created by cloning or deploying the same template, then all of these Windows virtual machines may end up using the same GUID on their dynamic volumes.
- File-level restore does not restore or browse symbolic links.
- File-level restore of Windows 8, Windows Server 2012 and Windows Server 2016 virtual machines is not supported on the following file systems:
 - Deduplicated NTFS
 - Resilient File System (ReFS)
 - EFI bootloder

FLR Agent for virtual machine file level restore

The VM Direct **FLR Agent** is required for file level restore operations and is installed automatically on the target virtual machine when you initiate a file level restore and provide the virtual machine credentials.

- NOTE:** The most up-to-date software compatibility information for PowerProtect Data Manager is provided in the E-Lab Navigator, available at <https://elabnavigator.emc.com/eln/modernHomeDataProtection>.

FLR Agent installation on Linux virtual machines

The **FLR Agent** installation on Linux virtual machines requires that you use the root account, or that you are a user in the operating system's local `sudoers` list. If credentials for any other user are provided for the target virtual machine, the **FLR Agent** installation fails, even if this user has privileges similar to a root user.

To allow a non-root user or group to perform the **FLR Agent** installation:

1. Provide sudo access to the following files at a minimum:
 - `rpm` command (SLES, RHEL, CentOS) and `dpkg` command (Debian/Ubuntu)
 - `/opt/emc/vproxyra/bin/postinstall.sh`
 - `/opt/emc/vproxyra/bin/preremove.sh`

Note the following additional requirements:

- The sudo user or group must be configured for no password prompt.
 - The sudo user or group must be provided with the `no_requiretty` option.
 - To browse files for a file level restore when you have user elevation enabled, you must have the appropriate authority in the guest virtual machine operating system. For example, you must be permitted to run `vflxbrowse` using `sudo` without prompting for a password.
 - To perform a file-level restore when you have user elevation enabled, you must have the appropriate authority. For example, you must be permitted to run `vflrcopy` using `sudo` without prompting for a password.
2. On the Linux system, create the file `/etc/sudoers.d/linuxuser`, where `linuxuser` is the Linux login user, and then add the following contents to this file.

On CentOS, Red Hat, SuSE, OpenSUSE, and Oracle Linux platforms:

```
username ALL=NOPASSWD: /usr/bin/sudo, /usr/bin/rpm, /opt/emc/
vproxyra/bin/postinstall.sh, /opt/emc/vproxyra/bin/preremove.sh, /opt/emc/vproxyra/bin/
vflrbrowse, /opt/emc/vproxyra/bin/vflrcopy
```

```
Defaults:username !requiretty
```

On Ubuntu platforms:

```
username ALL=NOPASSWD: /usr/bin/sudo, /usr/bin/dpkg, /opt/emc/
vproxyra/bin/postinstall.sh, /opt/emc/vproxyra/bin/preremove.sh, /opt/emc/vproxyra/bin/
vflrbrowse, /opt/emc/vproxyra/bin/vflrcopy
```

```
Defaults:username !requiretty
```

Once you complete the **FLR Agent** installation on the target virtual machine using the root user account or a sudouser with the minimum file access requirements, you can perform file level restore operations as a non-root user on supported Linux platforms. To determine which Linux platforms are supported, review the compatibility information at <https://elbnavigator.emc.com/elb/modernHomeDataProtection>.

FLR Agent installation on Windows virtual machines

FLR Agent installation on Windows virtual machines requires that you use administrative privileges. If the provided credentials for the target virtual machine do not have administrative privileges, the **FLR Agent** installation fails.

On Windows, to perform a file-level restore using a non-administrator user, ensure that the **FLR Agent** is already installed on the target machine using administrative privileges. Otherwise, ensure that an administrative user is specified, and click **OK**.

Installation of the **FLR Agent** on User Account Control (UAC) enabled Windows virtual machine requires you to either provide the credentials of the administrator user, or to disable UAC during the **FLR Agent** installation and then re-enable upon completion.

On Windows versions 7, 8, and 10, the administrator account is disabled by default. To enable the account, complete the following steps:

1. To activate the account, open a command prompt in administrative mode, and then type `net user administrator /active: yes`.
2. To set a password for the administrator account, go to **Control Panel > User Accounts** and select the **Advanced** tab. Initially, the account password is blank.
3. In the **User Accounts** pane, right-click the user and select **Properties**, and then clear the **Account is disabled** option.

To disable UAC during the **FLR Agent** installation and then re-enable on completion of the installation, complete the following steps:

1. Initiate a file-level restore to launch the **FLR Agent installation** window. The **FLR Agent** installation is automatically started during a mount operation if it is not already installed on the destination virtual machine.
2. In the **FLR Agent installation** window, select the **Keep VM Direct FLR on target virtual machine** option.
3. Open **regedit** and change the **EnableLUA** registry key value at `HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System` to `0x00000000`. By default, this is set to 1.
4. Proceed with the **FLR Agent** installation.
5. Open **regedit** and reset the **EnableLUA** registry key to the previous value to re-enable UAC.

Updating the Microsoft Application Agent and FLR Agent software

The **Microsoft Application Agent** and **FLR Agent** software required to perform SQL application-aware data protection and file-level restore operations will be automatically updated on the target virtual machine by the VM Direct appliance during the file-level restore operation. The VM Direct appliance detects the available software on the client and updates the Agent software with the new version of software from its repository. If the update does not occur automatically, contact a Dell EMC technical support professional for a procedure to update the VM Direct software repository with the latest version of the Agent software packages.

FLR-supported platform and OS versions for virtual machine restores

File-level restore is supported for the following platforms and operating system versions only.

Platforms/operating systems are qualified for file-level restore support using the default file system for these platforms:

NOTE: The most up-to-date software compatibility information for PowerProtect Data Manager is provided in the E-Lab Navigator, available at <https://elabnavigator.emc.com/eln/modernHomeDataProtection>.

- RedHat Enterprise Linux versions 6.x, 7.x, and 8.x
- SuSE Linux Enterprise Server versions 11.x and 12.x
- Debian version 9.1
- Ubuntu version 17.10
- CentOS version 7.2 and later
- Oracle Enterprise Linux version 7.2 and later
- Windows 7, 8, 10, Server 2008, 2012, 2016 (all 64-bit platforms and R2, where applicable), 2019 for FAT, and NTFS.

Ensure that the latest supported version of VMware Tools or open-vm-tools is installed on the guest operating system.

Support for Debian or Ubuntu operating system

vProxy file-level restore is supported on the Debian/Ubuntu operating system. To configure the Debian or Ubuntu guest operating system for file-level restore, perform the following steps.

About this task

NOTE: File-level restore is not supported on Debian ext4 file systems.

Steps

1. Log in to the system console as a non-root user.
2. Run the `sudo passwd root` command.
Enter the new password twice to set a password for the root account.
3. Run the `sudo passwd -u root` command to unlock the root account.
4. Specify the root user credentials in the **Dell EMC Data Protection Restore Client** and proceed to complete the file-level restore operation at least once.
While performing the file-level restore operation for the first time, remember to select **Keep FLR agent**.
5. After performing the above steps at least once, you can revert the root account to the locked state and use non-root account for future file-level restore requests. Non-root user can lock the root account with the `sudo passwd -l root` command.

Operating system utilities required for file-level restore

On Linux and Windows, the installed operating system must include several standard utilities in order to use file-level restore. Depending on the target operating system for restore and the types of disks or file systems in use, some of these standard utilities, however, may not be included.

The following utilities and programs may be required for performing file-level restore.

On Windows:

- msexec.exe
- diskpart.exe
- cmd.exe

On Linux:

- blkid
- udevadm
- readlink

- rpm
- bash

① **NOTE:** On Linux LVM, LVM2 rpm version 2.02.117 or later is required. Also, additional binaries required on Linux LVM include dmsetup, lvm, and vgimportclone.

PowerProtect Data Manager resource requirements in a VMware environment

Review the following minimum system requirements for PowerProtect Data Manager in a VMware environment (ESXi server).

- CPU—10 CPU cores
- Memory—18 GB RAM for PowerProtect Data Manager
- Seven disks with the following capacities:
 - Disk 1—100 GB
 - Disk 2—500 GB
 - Disks 3 and 4—10 GB each
 - Disks 5 through 7—5 GB each
- 1 GB network interface card (NIC)

① **NOTE:** If you plan to use Cloud DR, your system must also meet the following requirements:

- CPU—14 CPU cores
- Memory—22 GB

Software and hardware requirements

The following table lists the required components for PowerProtect Data Manager and the VM Direct protection engine.

Table 58. PowerProtect Data Manager and VM Direct engine requirements

Component	Requirements	Notes
PowerProtect Data Manager with the VM Direct engine	Version 19.9 or later.	
vCenter Server	vSphere and ESXi versions 6.5, 6.7, 7.0, 7.0 U1 and later.	<p>Refer to the VMware documentation ESXi 6.5 and later <i>minimum requirements for physical host requirements for the ESXi hosts.</i></p> <p>VMware has announced the end of general support for vSphere version 6.0. The Knowledge Base article at https://kb.vmware.com/s/article/66977 provides more information.</p> <p>Version 6.5 and later is required to perform Microsoft SQL Server application-aware protection. Also, file-level restore in the vSphere Client requires a minimum vCenter version 6.7 U1.</p> <p>Any new virtual machine protection policies use Transparent Snapshot Data Mover (TSDM) as the default protection mechanism instead of VADP, provided that the vCenter/ESXi Server that hosts the virtual machines is a minimum version of 7.0 U3 and the policy options selected for the virtual machine crash-consistent protection policy are supported by TSDM.</p>

Table 58. PowerProtect Data Manager and VM Direct engine requirements (continued)

Component	Requirements	Notes
VMware Tools	Version 10 or later.	Install VMware Tools on each virtual machine by using the vSphere Client . VMware Tools adds additional backup and recovery capabilities that quiesce certain processes on the guest operating system before backup. Version 10.1 and later is required to perform Microsoft SQL Server application-aware protection.
PowerProtect DD System models and software	<ul style="list-style-type: none"> All models of PowerProtect DD System in production are supported. DD Operating System (DDOS) version 6.2 or later and the PowerProtect DD Management Center (DDMC). 	Make note of the hosts writing backups to your DD systems.
Web browser	Google Chrome	The latest version of the Google Chrome browser is recommended to access the PowerProtect Data Manager UI.

Support for backup and restore of encrypted virtual machines

Backup and restore of encrypted virtual machines is supported in PowerProtect Data Manager, with the following limitations:

- Restoring encrypted virtual machines to a different vCenter Server is not supported. You must perform the restore to the original virtual machine or a new virtual machine in the same vCenter.
- Restoring an encrypted virtual machine backup to a new virtual machine in the original vCenter Server will restore the virtual machine disks (VMDKs) in clear text if the VMDKs are not encrypted. The article "Virtual Machine Encryption" at <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-E6C5CE29-CD1D-4555-859C-A0492E7CB45D.html> provides more information about manually changing the virtual machine policy to enable encryption of VMDKs.
- In order to use **Hot Add** transport mode, all VM proxies with access to the encrypted virtual machines datastore must be encrypted as well. For example, if encrypted virtual machines reside in an ESXi cluster, all VM proxies deployed on the cluster must also be encrypted.
- In order to backup and restore encrypted virtualization-based security (VBS) and virtual Trusted Platform Module 2.0 (vTPM) virtual machines, vCenter 7.0 U1 or later is required.

Transport mode considerations

Review the following information for recommendations and best practices when selecting a transport mode to use for virtual machine data protection operations and Tanzu Kubernetes guest cluster protection in PowerProtect Data Manager.

Hot Add transport mode recommended for large workloads

For workloads where full backups of large sized virtual machines or backups of virtual machines with a high data change rate are being performed, **Hot Add** transport mode provides improved performance over other modes. With **Hot Add** transport mode, a VM Direct Engine must be deployed on the same ESXi host or cluster that hosts the production virtual machines. During data protection operations, a VM Direct Engine capable of performing Hot Add backups is recommended. The following selection criteria is used during data protection operations:

- If a VM Direct Engine is configured in Hot Add only mode, then this engine is used to perform **Hot Add** virtual machine backups. If one or more virtual machines are busy, then the backup is queued until the virtual machine is available.
- If a virtual machine is in a cluster where the VM Direct Engine is not configured in **Hot Add** mode, or the VM Direct Engine with **Hot Add** mode configured is disabled or in a failed state, then PowerProtect Data Manager selects a VM Direct Engine within the cluster that can perform data protection operations in **NBD** mode. Any VM Direct Engine with **Hot Add** mode configured that is not in the cluster is not used.
- Any VM Direct Engine that is configured in **NBD** only mode, or in **Hot Add** mode with fallback to **NBD**, is used to perform **NBD** virtual machine backups. If every VM Direct Engine that is configured in **NBD** mode is busy, then the backup is queued until one of these engines is available.

- If there is no VM Direct Engine that is configured in **NBD** mode, or the VM Direct Engine with **NBD** mode configured is disabled or in a failed state, then the PowerProtect Data Manager embedded VM Direct engine is used to perform the **NBD** backup.

Other transport mode recommendations

Review the following additional transport mode recommendations:

- Use **Hot Add** mode for faster backups and restores and less exposure to network routing, firewall, and SSL certificate issues. To support **Hot Add** mode, deploy the VM Direct Engine on an ESXi host that has a path to the storage that holds the target virtual disks for backup.
 - ① **NOTE:** **Hot Add** mode requires VMware hardware version 7 or later. Ensure all virtual machines that you want to back up are using Virtual Machine hardware version 7 or later.
- In order for backup and recovery operations to use **Hot Add** mode on a VMware Virtual Volume (vVol) datastore, the VM Direct proxy should reside on the same vVol as the virtual machine.
- If you have vFlash-enabled disks and are using **Hot Add** transport mode, ensure that you configure the vFlash resource for the VM Direct host with sufficient resources (greater than or equal to the virtual machine resources), or migrate the VM Direct Engine to a host with vFlash already configured. Otherwise, backup of any vFlash-enabled disks fails with the error `VDDK Error: 13: You do not have access rights to this file` and the error on the vCenter server `The available virtual flash resource '0' MB ('0' bytes) is not sufficient for the requested operation.`
- For sites that contain many virtual machines that do not support **Hot Add** requirements, Network Block Device (**NBD**) transport mode is used. This mode can cause congestion on the ESXi host management network. Plan your backup network carefully for large scale **NBD** installs, for example, consider configuring one of the following options:
 - Setting up Management network redundancy.
 - Setting up backup network to ESXi for **NBD**.
 - Setting up storage heartbeats.

See <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/vmw-vsphere-high-availability-whitepaper.pdf> for more information.
- If performing **NBD** backups, ensure that your network has a bandwidth of 10 Gbps or higher.

Virtual disk types supported

When planning your protection policies, ensure that PowerProtect Data Manager supports the disk types that you use in the environment.

PowerProtect Data Manager does not support the following disk types:

- First Class Disks
- Independent (persistent and nonpersistent)
- RDM Independent - Virtual Compatibility Mode
- RDM Physical Compatibility Mode

Additionally, Dell EMC recommends to avoid deploying VMs with IDE virtual disks, which degrades backup performance. Use SCSI virtual disks instead whenever possible. Note that you cannot use **Hot Add** mode with IDE Virtual disks. Backup of IDE Virtual disks is performed using **NBD** mode.

Virtual machine data change rate

The data change rate is the percentage of a virtual machine's data that changes between backups.

Data change rates directly impact the number of VM Direct Engines required to successfully complete the backup of all required virtual machines within the backup window. A daily data change rate of 3-4% is typical in a vSphere environment. Higher data change rates will require either a longer window to complete the backup, additional VM Direct Engines, or both.

VM Direct Engine data ingestion rate

The VM Direct Engine data ingestion rate is another parameter that directly impacts the number of VM Direct Engines required to successfully complete the backup of all required virtual machines within the backup window.

By default, each VM Direct Engine processes approximately 500 GB to 1TB of data per hour, subject to the deduplication and read throughput on the primary stack. A number of additional factors, however, can impact the actual data ingestion rate, including the following:

- The protection storage system being used for data protection operations.
- The type of storage media used for VM Direct Engine storage.
- Your network and/or SAN infrastructure and connectivity speed.

If data ingestion rates at your site are typically lower or higher than 500 GB per hour, you can add or delete VM Direct Engines as needed. You can also shorten or lengthen the backup window. By default, each VM Direct Engine is configured to handle the optimal number of concurrent VMDK backup jobs. Configuring each VM Direct Engine to allow fewer concurrent backup jobs would typically require deploying additional VM DirectEngines, but can result in more evenly distributed backup jobs among each VM Direct Engine.

Full (Level-0) backups typically take longer and consume more VM Direct Engine resources. Therefore, large new virtual machine deployments can impact the ability to complete all required backups within the time specified for the backup window. In order to allow the system to perform these full backups without interruption, where possible ensure that you implement a phased approach for large new virtual machine deployments. If a phased deployment is not possible, and the full backups do not complete before timeout of the backup window, you can also enable automatic retry of failed backups. The section *Restart a job or task automatically* on page 148 provides instructions. It is recommended that an administrator user monitor such workloads to ensure that the system can handle these workloads when the demand on resources begins to decrease, and that the virtual machine backups then complete successfully.

VM Direct Engine limitations and unsupported features

Review the following limitations and unsupported features related to the VM Direct Engine.

Backup of individual folders within a virtual machine is not supported

PowerProtect Data Manager only supports image-level backup and disk-level backup. You cannot perform backups of individual folders within the virtual machine.

Backups fail for resource pools recreated with the same name as deleted pool

When you delete a resource pool in vCenter and then recreate a resource pool with the same name, backups fail. Re-configure the protection group with the newly created resource pool.

Datastore names cannot contain special characters

Using special characters in datastore names can cause problems with the VM Direct Engine, such as failed backups and restores. Special characters include the following: `% & * $ # @ ! \ / : * ? " < > | ,` and so on.

DD Boost over fibre channel not supported

PowerProtect Data Manager does not support DD Boost over fibre channel (DFC).

Error when changing configuration of many virtual machines at the same time

When configuring or unconfiguring many virtual machines (300 or more) in a protection policy, an error message might display indicating that the request is too large. You can click **OK** and proceed, but system performance will be impacted due to the size of the request. As a best practice, it is recommended to use protection rules to automatically determine which assets are assigned to protection policies when the assets are discovered.

Hot Add backups fail when datacenter names contain special characters

Virtual machine backups fail when the datacenter name contains special characters and the transport mode specified for VM Direct backups is **Hot Add only**. Avoid using special characters in the datacenter name, for example, "Datacenter_#2@3", or specify **Hotadd with fallback to Network Block Device** for the transport mode.

Hot Add backups fail when virtual machine protection policy configured with Virtual Flash Read Cache value

When using **Hot Add** transport mode for a virtual machine protection policy, the backup fails with the following error if configured with the Virtual Flash Read Cache (vFRC) value:

```
"Backup has FAILED. Failed to backup
virtual disk \"Hard disk <no.>\". Failed to initialize Block
Reader. Failed to open source VMDK \"<dataStore
name>/<VM Name.vmdk>\": VDDK Error: 13: You do not have
access rights to this file. (500)".
```

I/O contention when all Virtual Machines on a single data store

I/O contention may occur during snapshot creation and backup read operations when all Virtual Machines reside on a single datastore.

Limitations to SQL Server application consistent data protection

Review the SQL Server application-consistent protection support limitations in the section "Microsoft application agent for SQL Server application-aware protection" of the *PowerProtect Data Manager for Microsoft Application Agent SQL Server User Guide*.

Network configuration settings are not restored with virtual machine after recovery of a vApp backup

Network configuration settings are not backed up with the virtual machine as part of a vApp backup. As a result, when you restore a vApp backup, you must manually reconfigure the network settings.

NFC log level settings

To assist with I/O performance analysis, set the NFC log level in the VM Direct proxy configuration file to its highest value, for example, `vixDiskLib.nfc.LogLevel=4`. Setting the log level in the server for NFC asynchronous I/O is not required. You can then run the VDDK sample code and evaluate I/O performance by examining the `vddk.log` and the `vpva` log file.

NOTE: Virtual Machines with very high I/O might stall during consolidation due to the ESXi forced operation called synchronous consolidate. Plan your backups of such Virtual Machines according to the amount of workload on the Virtual Machine.

Protection fails for virtual machine name containing { or }

A PowerProtect Data Manager virtual machine protection policy fails to back up virtual machines that contain the special characters { or } in the name. This limitation exists with vSphere versions previous to 6.7. If you do not have vSphere 6.7 or later installed, avoid using these two characters in virtual machine names.

SAN transport mode not supported

PowerProtect Data Manager supports only the Hot Add and NBD transport modes. The Hot Add mode is the default transport mode. For a protection policy, you can specify to use only Hot Add mode, only NBD mode, or Hot Add mode with fallback to NBD if Hot Add is not available.

Specify NBD for datastores if VM Direct should use NBD mode only

For a VM Direct Engine that will only use NBD transport mode, you must also specify the datastores for which you want the proxy to perform only NBD backups to ensure that any backups of virtual machines running on these datastores are always performed using NBD mode. This also ensures that the same NBD-only proxies are never used for backups of virtual machines residing on any other datastores.

Thin provisioning not preserved during NFS datastore recovery

When backing up thin-provisioned virtual machines or disks for virtual machines on NFS datastores, an NFS datastore recovery does not preserve thin provisioning. VMware knowledge base article 2137818 at <https://kb.vmware.com/kb/2137818> provides more information.

Virtual machine alert "VM MAC conflict" may appear after successful recovery of virtual machine

After performing a successful recovery of a virtual machine through vCenter version 5, an alert may appear indicating a "VM MAC conflict" for the recovered virtual machine, even though the new virtual machine will have a different and unique MAC address. You must manually acknowledge the alert or clear the alert after resolving the MAC address conflict. Note that this alert can be triggered even when the MAC address conflict is resolved.

The VMware release notes at <https://docs.vmware.com/en/VMware-vSphere/6.0/rn/vsphere-vcenter-server-60u2-release-notes.html> provide more information.

VM Direct Engine configuration settings cannot be modified after adding the VM Direct Engine

After adding a VM Direct Engine, the only field you can modify is the **Transport Mode**. Any other configuration changes require you to delete and then re-add the VM Direct Engine. Additional VM Direct actions on page 68 provides more information.

VM Direct Engine configured with dual stack is not supported

The VM Direct Engine does not support dual stack (IPv4 and IPv6) addressing. If you want to run backups and restores using the VM Direct Engine, use IPv4 only addressing.

VMware Distributed Resource Scheduler cluster support limitations

The PowerProtect Data Manager server is supported in a VMware Distributed Resource Scheduler (DRS) cluster, with the following considerations:

- During backup of a virtual machine, `host-vmotion` or `storage-vmotion` is not permitted on the virtual machine. The option to migrate will be disabled in the **vSphere Client** UI.
- If the VM Direct proxy is in use for a backup or restore with **Hot Add** disks attached, then `storage-vmotion` of the vProxy is not permitted during these operations.

VMware limitations by vSphere version

VMware limitations for vSphere 6.0 and later versions are available at <https://configmax.vmware.com/home>. For vSphere 5.5, go to <https://www.vmware.com/pdf/vsphere5/r55/vsphere-55-configuration-maximums.pdf>.

VMware snapshot for backup is not supported for independent disks

When using independent disks you cannot perform VMware snapshot for backup.

VM Direct Engine performance and scalability

The VM Direct Engine performance and scalability depends on several factors, including the number of vCenter Servers and proxies and the number of concurrent virtual machine backups. The following table provides information on these scalability factors and maximum recommendations, in addition to concurrence recommendations for sessions created from backups using the VM Direct Engine.

The count of sessions is driven by the number of proxies and backups running through this server.

Table 59. Performance and scalability factors

Component	Maximum limit	Recommended count	Notes
Number of concurrent NBD + Preferred Hot Add backups per ESXi host	48		Ensure that your network has a bandwidth of 10 Gbps or higher. VMware uses Network File Copy (NFC) protocol to read VMDK using NBD transport mode. You need one VMware NFC connection for each VMDK file being backed up. The VMware Documentation provides more information on vCenter NFC session connection limits.
Concurrent VMDK backups per vCenter Server		180	Can be achieved with a combination of the number of proxies multiplied by the number of configured Hot Add sessions per VM Direct Engine.
Number of proxies per vCenter Server	25	7	A limit of 25 concurrent backup and recovery sessions.
Number of files/directories per file level recovery	200,000		File-level restore is recommended for quickly recovering a small set of files. Image-level or VMDK-level recoveries are optimized and recommended for recovering a large set of files/folders.

When you reach the limit for concurrent backup sessions, a warning message displays. The remaining sessions will be queued. You can adjust the session limits by modifying the `MAX_VC_BACKUP_SESSIONS` and `MAX_NBD_BACKUP_SESSIONS` variables in the environment file, according to the recommendations. The Knowledge Base article 543253 at <https://support.emc.com/kb/543253> provides more information.

Table 60. Proxy session limits by proxy type

Component	Total number of sessions (backup and recovery) maximum	Notes
Added (External) VM Direct Engine	25	
Embedded VM Direct engine NOTE: The embedded VM Direct engine is pre-bundled with the PowerProtect Data Manager software.	4	The embedded VM Direct engine is only used as a fallback when all other proxies are disabled or in Failed state.

VM Direct Engine selection with virtual networks (VLANs)

PowerProtect Data Manager typically selects a VM Direct Engine by accounting for availability, transport mode settings, and engine load. This selection optimizes data throughput.

When you configure virtual networks for PowerProtect Data Manager and VM Direct Engine to isolate backup traffic, you can define routes to the protection storage system interface for each virtual network. The routes that you configure can influence VM Direct Engine selection. PowerProtect Data Manager ensures that the selected engine has a network interface that can send traffic for a specific virtual network to the protection storage system.

Best practices for vCenter Server backup and restore

Review the following recommendations and best practices when planning a vCenter Server backup and restore.

NOTE: Backups will not save Distributed switch configurations.

- It is recommended to schedule the backup of the vCenter Server when the load on the vCenter Server is low, such as during off-hours, to minimize the impact of vCenter virtual machine snapshot creation and snapshot commit processing overhead.
- Ensure that there are no underlying storage problems that might result in long stun times.
- Keep the vCenter virtual machine and all of its component virtual machines in one single isolated protection policy. The protection policy should not be shared with any other virtual machines. This is to ensure that the backup times of all vCenter Server component virtual machines are as close to each other as possible.
- Ensure that the backup start time of the vCenter Server does not overlap with any operations for other protected virtual machines being managed by this vCenter so that there is no impact on other protected virtual machines during snapshot creation and snapshot commit of the vCenter virtual machine.
- If the vCenter Server and Platform Services Controller instances fail at the same time, you must first restore the Platform Services Controller and then the vCenter Server instances.

Changing the vCenter server FQDN

If you change the fully qualified domain name (FQDN) of the vCenter server, PowerProtect Data Manager must be reconfigured to accommodate this change without any issues.

When the FQDN of the vCenter server changes, so does its SSL certificate. In order to continue to administer the vCenter server and maintain uninterrupted protection of its assets, the new certificate must be imported into the PowerProtect Data Manager trust store.

Change the vCenter server FQDN

When the FQDN of the vCenter server changes, its new SSL certificate must be imported into the PowerProtect Data Manager trust store.

About this task

This procedure uses REST API commands that are run on the PowerProtect Data Manager server.

NOTE: In the following steps, replace **192.168.1.204** with the IP address of the PowerProtect Data Manager server and **a022-renamed-ppdm.vmware.com** with the new FQDN of the vCenter server.

Steps

1. Get the current information from the vCenter server, and make a note of the value of *id*, which corresponds to the new FQDN certificate:

```
GET https://192.168.1.204:8443/api/v2/certificates?host=a022-renamed-ppdm.vmware.com&port=443&type=Host
```

For example, the output might look like this:

```
fingerprint: "43FF8FBA62D1DD68E630AE9DB8BA7DF21549CE39"  
host: " a022-renamed-ppdm.vmware.com"
```

```
id: "dnN1bnRlcil2bWRtLTAA0LmFzbC5sYWluZW1jLmNvbTo0NDM6aG9zdA=="
issuerName: "OU=VMware Engineering, O= a022-renamed-ppdm.vmware.com, ST=California, C=US, DC=local, DC=vsphere, CN=CA"
notValidAfter: "Mon Mar 11 17:39:09 PDT 2030"
notValidBefore: "Mon Mar 16 17:39:09 PDT 2020"
port: "443"
state: "UNKNOWN"
subjectName: "C=US, CN=vcenter-vmdn-04.asl.lab.emc.com"
type: "HOST"
```

2. Import the new certificate into the PowerProtect Data Manager trust store:

```
PUT https://192.168.1.204:8443/api/v2/certificates/{newCertID}
```

Replace *{newCertID}* with the value of *id* displayed in step 1. Only use the text that was displayed between the quotation marks.

3. Get the ID of the vCenter server:

```
GET https://192.168.1.204:8443/api/v2/inventory-sources/
```

All vCenter servers that are configured in PowerProtect Data Manager are displayed.

For example, the output might look like this:

```
"id": "6ffdb6e9-b864-56f4-8ec8-felc214c6fef",
  "name": "VC",
  "version": "7.0.2",
  "type": "VCENTER",
  "lastDiscovered": "2021-08-10T07:03:41.624Z",
  "lastDiscoveryResult": {
    "status": "OK",
```

4. Record the new FQDN of the vCenter server in PowerProtect Data Manager:

```
PUT https://192.168.1.204:8443/api/v2/inventory-sources/{vCenter-id}
```

Replace *{vCenter-id}* with the value of *id* displayed for the vCenter in step 3. Only use the text that was displayed between the quotation marks.

5. Get the current list of certificates:

```
GET https://192.168.1.204:8443/api/v2/certificates
```

Both the old and new FQDN certificates are displayed. There might also be additional certificates displayed.

6. Search the certificate entries displayed in step 5, and locate the entry where the value of *host* matches the old FQDN of the vCenter server. Make a note of the corresponding *id* value.

7. Delete the old certificate from the PowerProtect Data Manager :

```
DELETE https://192.168.1.204:8443/api/v2/certificates/{oldCertID}
```

Replace *{oldCertID}* with the value of *id* noted in step 6. Only use the text that was displayed between the quotation marks.

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 on page 73 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.

Replacing security certificates

You can replace the default self-signed security certificates for the PowerProtect Data Manager UI, or replace changed or expired security certificates on an external server.

The *PowerProtect Data Manager Security Configuration Guide* provides more information.

Replacing the self-signed security certificates

If you want to use certificates for the PowerProtect Data Manager UI that are signed by a certificate authority (CA) of your choice, you can replace them.

The *PowerProtect Data Manager Security Configuration Guide* provides more information.

Replace expired or changed certificates on an external server

Use this procedure to replace certificates on an external server (for example, a DD, LDAPS, or vCenter server) that have expired or changed. Only the Administrator role can replace certificates.

About this task

If a certificate on the external server has expired or been changed, connection to the server fails with the following error:

```
javax.net.ssl.SSLHandshakeException: sun.security.validator.ValidatorException: PKIX
```

Perform the following steps using cURL or any REST API client, such as Postman.

Steps

1. Log in to the external server as an administrator:

```
POST https://server_hostname:REST port number/api/v2/login
```

Provide the following request payload in JSON format:

```
{
  "username": "username",
  "password": "password"
}
```

where *username* is a user with the Administrator role and *password* is the password for this user.

NOTE: Add the following header key with your REST call request:

```
'Content-type: application/json'
```

The response returns the following information:

```
{
  "access_token":
  "token_type":
  "expires_in":
  "jti":
  "scope":
  "refresh_token":
}
```

Copy the **access_token** value from the response above. This value will be required in the header key **Authorization** for all the REST calls in subsequent steps.

2. On the REST API client, run the following to obtain the old or expired external server certificate:

```
GET https://server_hostname:REST port number/api/v2/certificates
```

NOTE: Add the following header key with your REST call request:

```
'Authorization: access_token_value'
```

The response returns a list of certificate entries, each containing the following information:

```
{
  "id":
  "host":
  "port":
  "notValidBefore":
  "notValidAfter":
  "fingerprint":
  "subjectName":
  "issuerName":
  "state":
  "type":
}
```

NOTE: Make note of the **host**, **port** and **type** of each certificate, as this information will be required in Step 4. If you supply incorrect information in Step 4, requests that use these external hosts might fail.

3. On the REST API client, delete the old or expired external server certificate from the PowerProtect Data Manager datastore, using the ID obtained from the response in step 2:

DELETE `https://server hostname:REST port number/api/v2/certificates/id`

NOTE: Add the following header key with your REST call request:

```
'Authorization: access_token_value'
```

Ensure that you delete only the external server certificate that you want to remove.

4. On the REST API client, obtain the new certificate from the external server, using the host, port, and type obtained from the response in step 2:

GET `https://server hostname:REST port number/api/v2/certificates?host=host&port=port&type=type`

NOTE: Add the following header key with your REST call request:

```
'Authorization: access_token_value'
```

The response returns the following information:

```
{
  "id":
  "host":
  "port":
  "notValidBefore":
  "notValidAfter":
  "fingerprint":
  "subjectName":
  "issuerName":
  "state": "UNKNOWN",
  "type":
}
```

5. On the REST API client, accept the new certificate, using the ID obtained in the response from step 4:

PUT `https://server hostname:REST port number/api/v2/certificates/id`

NOTE: Add the following header key with your REST call request:

```
'Authorization: access_token_value'
```

Also, copy the response payload from step 4 in JSON format and change the state from "UNKNOWN" to "ACCEPTED".

6. On the REST API client, verify that the new certificate has been accepted, using the ID obtained in the response from step 4:

GET `https://server hostname:REST port number/api/v2/certificates/id`

NOTE: Add the following header key with your REST call request:

```
'Authorization: access_token_value'
```

If the certificate was accepted, the response returns the following information:

```

[[
  "id":
  "host":
  "port":
  "notValidBefore":
  "notValidAfter":
  "fingerprint":
  "subjectName":
  "issuerName":
  "state": "ACCEPTED",
  "type":
]]

```

Restarting PowerProtect Data Manager

When a PowerProtect Data Manager restart is required, Dell Technologies recommends 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`.

Scalability limits for vCenter Server, VM Direct Engine and DD systems

The following limits have been tested successfully with PowerProtect Data Manager for the vCenter Server, VM Direct Engine, and DD systems.

NOTE: These numbers are not maximum or hard limits, but should be considered when scaling your environment.

Table 61. Scalability limits

Component	Tested limits
Number of vCenter Servers supported with a single PowerProtect Data Manager server	12 NOTE: The vCenter server limit is subject to the VM Direct Engine overall limit of 40 and the per vCenter server limit of 25. For example, using the maximum tested number of vCenter servers of 12, you could add an average of 3 VM Direct Engines per vCenter server.
Number of external VM Direct Engines supported with a single PowerProtect Data Manager server	40 NOTE: This number was tested across 10 vCenter servers. For example, 4 VM Direct Engines per vCenter server.
Number of DD systems supported per PowerProtect Data Manager server	10
Network latency between the PowerProtect Data Manager server and VM Direct Engines	200 ms
Network latency between the PowerProtect Data Manager server and the DD systems	200 ms
Number of virtual machines per PowerProtect Data Manager server	10,000

Troubleshooting network setup issues

vCenter registration and proxy deployment fails if the PowerProtect Data Manager server is deployed in the same private network as the internal Docker network.

PowerProtect Data Manager uses an internal private Docker network. If the PowerProtect Data Manager server is deployed in the same private network as the internal Docker network, or if some data sources have already been deployed within the private network, PowerProtect Data Manager fails to protect the data sources.

To resolve this issue, deploy the PowerProtect Data Manager server and other data sources in a different network. If you cannot modify the deployed network, run a script tool within PowerProtect Data Manager to switch the private Docker network to a different network.

To switch the private Docker network to a different network:

1. Connect to the PowerProtect Data Manager console and change to the root user.
2. Modify the Docker network by running the following command:

```
/usr/local/brs/puppet/scripts/docker_network_switch.sh subnet gateway
```

Where:

- `subnet` describes the new network in the format `172.25.0.0/24`
- `gateway` is the gateway for the private network. For example: `172.25.0.1`

Ensure that you specify a subnet and gateway that is not in use.

Troubleshooting PowerProtect agent service installations

A PowerProtect agent service installation might fail with the following error message:

```
Service 'PowerProtect Agent Service' (AgentService) could not be installed. Verify that you have sufficient privileges to install system services.
```

Possible causes of the installation failure are as follows:

- The installation was attempted on a passive node of a Failover Cluster Instance (FCI).
- The installation was canceled and a rollback left some stale entries of PowerProtect agent services.

As a workaround, clean up the PowerProtect agent service entries, and retry the installation.

Troubleshooting PowerProtect agent service operations

When investigating issues with the PowerProtect agent service, you might need to troubleshoot its operations:

Troubleshoot the PowerProtect agent service operations

To troubleshoot the agent service operations, you can check the 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.

About this task

To modify the log level and retention of temporary files, you can perform the following steps.

Steps

1. Stop the agent service by using the appropriate procedure from the preceding topic.
2. Open the `config.yml` file in an editor.
3. Modify the log-level settings in the following parameters, as required:

NOTE: These parameters are listed in order of decreasing number of messages in the debug information output. The default log-level is `INFO`.

- `DEBUG`
- `INFO`
- `WARNING`
- `ERROR`
- `CRITICAL`

4. 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.

5. Start the agent service by using the appropriate procedure from the preceding topic.

Troubleshooting PowerProtect Data Manager software updates

Review the following information related to updating the PowerProtect Data Manager software.

Mounting a read-only file system results in a failed update

If you mount a read-only file system under the `/home/admin` or `/home/sysadmin` directories on the PowerProtect Data Manager node, the update cannot complete successfully. Ensure that you remove read-only file system mounts before updating PowerProtect Data Manager.

Managing certificates after updating PowerProtect Data Manager from versions earlier than 19.1

Use this procedure to ensure that certificates existing on the pre-update system also exist on the post-update system.

Prerequisites

Ensure that you update any expired certificates on external systems to valid certificates.

Steps

1. Connect to the PowerProtect Data Manager console as an admin user.
2. Run the upgrade command:

```
/usr/local/brs/lib/secretmgr/bin/secretmgr-tls-upgrade
```

The system displays the external system certificates.

3. Verify each certificate as trusted or untrusted: At the prompt for each certificate, type `Y` to accept. Any other character rejects the certificate. Expired certificates are automatically rejected.

Troubleshooting storage units

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.

Troubleshooting virtual machine backup issues

This section provides information about issues related to virtual machine backup operations with the VM Direct protection engine.

Backup completes with a non-quieted snapshot warning

A virtual-machine backup completes, but with a warning that a non-quieted snapshot was used. Although most data will be protected, using a non-quieted snapshot can result in some data being out of date or missing altogether.

The following warning is seen after a backup completes:

```
Warnings occurred during snapshot creation. Non-quieted snapshot was used,
quieted snapshot was unsuccessful. Unable to create quieted snapshot: An error
occurred while quieting the virtual machine. See the virtual machine's event log for details.
```

This can happen with backups of both Windows and Linux virtual machines. Refer to the following procedures for common methods of resolving the issue.

Troubleshooting non-quieted Windows snapshots

There is a common method of resolving this issue on Windows.

Steps

1. Confirm that the virtual machine has VMware Tools 10.1.0 or higher installed. If the virtual machine does not have VMware Tools 10.1.0 or higher installed, then install it.
2. Confirm that the *VMware Snapshot Provider* service is installed on the virtual machine. If the *VMware Snapshot Provider* service is not installed, then install it by reinstalling VMware Tools.

NOTE: Antivirus software might interfere with the installation of this service. If it is still not installed after reinstalling VMware Tools, then temporarily disable any antivirus software and reinstall VMware Tools again.

Troubleshooting non-quieted Linux snapshots

There is a common method of resolving this issue on Linux.

Steps

1. At a shell prompt of the virtual machine, run the command `cat /etc/vmware-tools/tools.conf`, and look for the value of `enableSyncDriver`:

```
[root]# cat /etc/vmware-tools/tools.conf
[vmbackup]
enableSyncDriver = false
```

2. If the value of `enableSyncDriver` is `false`, perform the following steps:
 - a. Edit `/etc/vmware-tools/tools.conf`, and change `enableSyncDriver = false` to `enableSyncDriver = true`.
 - b. At the shell prompt, run the command `systemctl restart vmtoolsd.service`.

Backup fails when names include special characters

When spaces or special characters are included in the virtual machine name, datastore, folder, or datacenter names, the .vmx file is not included in the backup.

The VM Direct appliance does not back up objects that include the following special characters (format: character/escape sequence):

- & %26
- + %2B
- / %2F
- = %3D
- ? %3F
- % %25
- \ %5C
- ~ %7E
-] %5D

Deleting vCenter asset sources or moving ESXi to another vCenter

When you delete a vCenter Server asset source from PowerProtect Data Manager without removing any vProxy/Search Nodes that the vCenter is hosting, the Nodes will become non-operational and move into Failed status upon the next health check. As a result, PowerProtect Data Manager updates will fail. This issue also occurs when you move the ESXi hosting the vProxy/Search Nodes from one vCenter to another vCenter.

To correct this issue, you can perform one of the following actions:

- Manually delete the vProxy/Search Nodes. The section *Delete vProxy/Search Nodes when vCenter Server asset source is no longer required* on page 266 provides the required steps.
- Return the vProxy/Search Nodes to an Operational/Ready state using the `vproxymgmt` and `infranodemgmt` tools. Choose this action if you want to add the vCenter again, or you want to add the vCenter that the ESXi has been moved to. The section *Return vProxy/Search Nodes to operational state when re-adding vCenter* on page 267 provides the required steps.

Delete vProxy/Search Nodes when vCenter Server asset source is no longer required

Perform the following procedure when you delete a vCenter server as an asset source in PowerProtect Data Manager and you will not be re-adding the vCenter:

About this task

NOTE: Manual cleanup of the virtual machine for the vProxy/Search Node has to be performed from the vCenter Server.

Steps

1. Run the following command to source the environment file.
`source /opt/emc/vmdirect/unit/vmdirect.env`
2. For vProxy removal:
 - a. Obtain the list of vProxies that require removal by running `/opt/emc/vmdirect/bin/vproxymgmt get`
 - b. Make note of the ID of any vProxy that needs to be deleted.
 - c. Use the `vproxymgmt` tool to delete vProxies by running `/opt/emc/vmdirect/bin/vproxymgmt delete -vproxy_id ProxyID`
3. For Search Node removal:
 - a. Obtain the list of Search Nodes that require removal by running `/opt/emc/vmdirect/bin/infranodemgmt get`
 - b. Make note of the ID of any Search Node that needs to be deleted.
 - c. Use the `infranodemgmt` tool to delete Search Nodes by running `/opt/emc/vmdirect/bin/infranodemgmt delete -node_id NodeID`
4. In the PowerProtect Data Manager UI, ensure that any sessions have been removed for both the vProxy/Search Node.

Return vProxy/Search Nodes to operational state when re-adding vCenter

When you want to re-add a vCenter that you deleted from PowerProtect Data Manager, or you want to add a vCenter that an ESXi has been moved to, perform the following procedure in order to return the vProxy/Search Nodes to an Operational/Ready state.

Steps

1. Re-add the deleted vCenter as an asset source in the PowerProtect Data Manager UI, or note the name of the new vCenter where the ESXi has been moved.
2. Run the following command to source the environment file.
`source /opt/emc/vmdirect/unit/vmdirect.env`
3. For vProxy updates:
 - a. Obtain the list of vProxies that require updating by running `/opt/emc/vmdirect/bin/vproxymgmt get`
 - b. Make note of the ID of any vProxy that needs to be updated.
 - c. Use the `vproxymgmt` tool to update the vCenter name by running `/opt/emc/vmdirect/bin/vproxymgmt modify -vcenter_hostname vCenter-FQDN -vproxy_id ProxyID`
4. For Search Node updates:
 - a. Obtain the list of Search Nodes that require updating by running `/opt/emc/vmdirect/bin/infranodemgmt get`
 - b. Make note of the ID of any Search Node that needs to be updated.
 - c. Use the `infranodemgmt` tool to update the vCenter name by running `/opt/emc/vmdirect/bin/infranodemgmt modify -vcenter_hostname vCenter-FQDN -node_id NodeID`
5. In the PowerProtect Data Manager UI, ensure that any sessions for the vProxy/Search Node and Cluster have changed to Operational/Ready state.

Failed to lock Virtual Machine for backup: Another EMC vProxy operation 'Backup' is active on VM

This error message appears when a backup fails for a virtual machine or when a previous backup of the virtual machine was abruptly ended and the VM annotation string was not cleared.

To resolve this issue, clear the annotation string value for the virtual machine.

1. Connect to the vCenter server and navigate **Home > Inventory > Hosts and Clusters**.
2. Select the virtual machine, and then select the **Summary** tab.
3. Clear the value that appears in the **EMC Proxy Session** field.

Lock placed on virtual machine during backup and recovery operations continues for 24 hours if VM Direct appliance fails

During VM Direct backup and recovery operations, a lock is placed on the virtual machine. If a VM Direct appliance failure occurs during one of these sessions, the lock is extended to a period of 24 hours, during which full backups and transaction log backups will fail with the following error until the lock is manually released:

```
Cannot lock VM 'W2K8R2-SQL-2014' (vm-522): Another EMC vProxy operation 'Backup' is active on VM vm-522.
```

Workaround

To manually release the lock on the virtual machine:

1. Open the **vSphere Web Client**.
2. Select the virtual machine and select **Summary**.
3. Select **Custom attribute** and click **Edit**.
4. Remove the attribute **EMC vProxy Session**.

Managing command execution for VM Proxy Agent operations on Linux

The VM Proxy Agent automatically creates a PAM service file named `vproxyra` in the `/etc/pam.d/system` directory, if the file does not already exist.

This file, which enables you to manage command execution through the VM Proxy Agent, is modeled on the corresponding `vmtoolsd` file. The settings in this file permit command execution by any user who is able to perform VM Direct operations on the guest virtual machine. A system administrator can further modify this file to specify which users can perform VM Direct operations, for example, file-level restore and SQL application-aware protection. For more information on the configuration of PAM service files, see the system documentation for your specific guest virtual machine operating system.

PowerProtect plug-in and portlet for vSphere display errors after replacing security certificates

After you replace the default self-signed security certificates, you may see errors in the vSphere client PowerProtect portlet when you select virtual machines:

- Service Unavailable: Please contact your administrator.
- No healthy upstream.

Reinstall the PowerProtect plug-in to apply the new certificates. The *PowerProtect Data Manager Security Configuration Guide* provides more information.

SQL databases skipped during virtual machine transaction log backup

When an advanced application-consistent policy is enabled with transaction log backup, the `mvagent_appbackup.exe` program evaluates databases to determine if transaction log backup is appropriate.

If transaction log backup is not appropriate for a database, the database will automatically be skipped. Databases are skipped for the reasons outlined in the following table.

Table 62. SQL Skipped Database Cases and Descriptions

Case	Description
Database has been restored	When a database has been restored, this database will be skipped during transaction log backup because there is no backup promotion.
System Database	System databases are automatically skipped for transaction log backup.
Database State	Database is not in a state that allows backup. For example, the database is in the NORECOVERY state.
Recovery Model	Database is in SIMPLE recovery model, which does not support transaction log backup
Other Backup Product	Most recent backup for the database was performed by a different backup product.
New Database	Database was created after most recent full backup.
Backup Failure	Database was in state to allow backup, backup was attempted, but backup failed.

All skipped databases will be backed up as part of the next full backup. Also, a skipped database will not result in `mvagent_appbackup.exe` failure. The only instance in which `mvagent_appbackup.exe` would potentially fail is if all databases failed to back up.

The `mvagent_appbackup.exe` program generates a history report of the databases, if the database backup status was success/skipped/failed, and a reason if they were skipped or failed if applicable. This history report is visible in the action logs for the VM Direct Engine, which are available as part of the `appbackup` logs.

NOTE: For SQL virtual machine application-consistent data protection, the SQL and operating system versions follow the support matrix available at <http://compatibilityguide.emc.com:8080/CompGuideApp/>.

SQL Server application-aware backup displays an error about `disk.EnableUUID` variable

Issue

A SQL Server application-aware virtual machine backup succeeds but displays the following error when the `disk.EnableUUID` variable for the virtual machine is set to `TRUE`:

```
VM '<asset_name>' configuration parameter 'disk.EnableUUID' cannot be evaluated. Map item 'disk.EnableUUID' not found. (1071)
```

Workaround

After you set the `disk.EnableUUID` variable to `TRUE`, reboot the virtual machine.

SQL Server application-consistent backups fail with error "Unable to find VSS metadata files in directory"

SQL Server application-consistent virtual machine backups might fail with the following error when the `disk.EnableUUID` variable for the virtual machine is set to `False`.

```
Unable to find VSS metadata files in directory C:\Program Files\DPSAPPS\MSVMAPPAGENT\tmp\VSSMetadata.xxxx.
```

To resolve this issue, ensure that the `disk.EnableUUID` variable for the virtual machines included in an SQL Server application-consistent backup is set to `True`.

Trailing spaces not supported in SQL database names

Due to a VSS limitation, you cannot use trailing spaces within the names of SQL databases protected by an application-consistent data protection policy.

VMware knowledge base articles and product documentation

Additional VMware troubleshooting information is available at the VMware Knowledge Base and VMware Documentation websites.

Troubleshooting virtual machine restore issues

The following topics provide information on troubleshooting virtual machine restore failures.

Virtual machine restores fail when `vProxyd` or `vrecoverd` disruption occurs

A virtual machine restore hangs and VPOD will not be able to reconnect to the restore session when the following scenarios occur:

- A disruption to the `vrecoverd` process on any external VM Direct engine.
- A disruption to the `vProxyd` process during a **Restore to Original Folder and Overwrite Original Files** or **Create and Restore to New VM** operation that uses Transparent Snapshot Data Mover (TSDM) as the protection mechanism.

After several retry attempts, VPOD marks the restore session as "Failed" and releases the `vProxy` associated with the restore.

If this failure occurs during a **Create and Restore to New VM**, you can delete the new virtual machine and restart the restore operation.

If this failure occurs during a **Restore to Original Folder and Overwrite Original Files**, you must remove the vProxy lock on the virtual machine from the vCenter, and then retry the restore operation. In the **vSphere Client**, the vProxy lock appears as a custom attribute with the name `Dell EMC vProxy Session`.

NOTE: If this attribute contains any value after a vProxyd process failure, backup and restore operations on this virtual machine cannot be performed. Clean up of this attribute and then running a successful restore operation is a requirement in order to avoid any potential data loss or corruption of the virtual machine, otherwise subsequent backups might also contain corrupted data.

DD NFS share not removed after restore to original

The NFS share might not be removed after a successful virtual machine restore to original. When this occurs, the restore hangs and the following NFS clients appear enabled in the DD system.



Figure 16. DD NFS clients still enabled after restore

If you encounter this issue, you can wait 24 hours for PowerProtect Data Manager to clean up the DD NFS shares, or you can stop the restore and clean up the DD NFS clients manually by performing the following steps:

1. Restart the VMDM service by typing `/usr/local/brs/lib/vmdm/bin/vmdm restart`.
2. Clean up DD NFS clients by typing `nfs del <Path> <Client>`.
3. In the vSphere Client's **Configuration** tab, manually unmount the `EMC-vProxy-vm-qa-xxxxxx` DDNFS datastore that is mounted on the ESXi host.

IP address change required after successful image-level restore to a new virtual machine

After performing a successful image-level restore to a new virtual machine, ensure that you change the IP address immediately in order to avoid IP conflicts with the original virtual machine. If you do not change the IP to a unique value, subsequent data protection operations might fail on the restored virtual machines, even if that virtual machine's network interfaces are disconnected.

Virtual machine protection copy does not display under available copies

If a virtual machine protection copy does not display under the available copies in PowerProtect Data Manager, verify the following:

- Ensure that protection of the virtual machine completed successfully.
- Check that the desired copy has not expired according to the PowerProtect Data Manager protection policy.

Virtual machine restore fails with name resolution error

A virtual machine restore might fail with the following error due to network issues between protection storage and PowerProtect Data Manager or the vCenter/ESXi:

```
com.emc.brs.vmdm.http.HttpsConnector - null: Temporary failure in name resolution  
java.net.UnknownHostException : null: Temporary failure in name resolution
```

Ensure that you have proper name resolution between protection storage and PowerProtect Data Manager /vCenter/ESX.

Virtual machine restore fails when the previous restore of this virtual machine is in progress or did not complete

A virtual machine restore fails with the following error if the previous restore operation for the same virtual machine is still in progress or did not complete successfully:

```
Error : There is another running restore operation that conflicts with this request.
```

If the previous restore operation for this virtual machine is still in progress, monitor the progress in PowerProtect Data Manager until the restore completes. If the virtual machine restore is complete but the task stops responding, then you must manually cancel the restore in PowerProtect Data Manager by restarting the VMDM service. You can restart the VMDM service by typing `/usr/local/brs/lib/vdm/bin/vdm restart`.

Virtual machine restore fails with error due to VM Direct corruption

A virtual machine restore might fail with the following error due to corruption of the VM Direct Engine that is running in PowerProtect Data Manager:

```
com.emc.dpsg.vproxy.client.VProxyManager - Error(createSession):  
javax.net.ssl.SSLException:  
Unrecognized SSL message, plaintext connection
```

Ensure that the `vproxyd` service is running in PowerProtect Data Manager by typing the following command.

```
ps xa | grep vproxy
```

Ensure that the `vproxy rpm` is installed as expected in PowerProtect Data Manager by typing the following command.

```
rpm -qa | grep vProxy
```

When logged in as the root user, restart the `vproxyd` service on PowerProtect Data Manager by typing the following command.

```
systemctl restart vproxyd
```

Virtual machine restore fails with error "User UserEARA does not have proper privileges"

A virtual machine restore fails with the error "User UserEARA does not have proper privileges" when the user does not have adequate privileges to perform the restore operation.

Ensure that the PowerProtect Data Manager user performing the restore belongs to System Tenant and has the Administrator or Restore Administrator role.

Troubleshooting instant access restore failures

An instant access restore consists of two stages. First, a virtual machine is made available in the UI as an instant access virtual machine without moving the virtual machine to permanent storage. Second, storage vMotion is initiated to migrate the virtual machine to permanent storage.

If at any point during the migration a restore failure occurs, the instant access session is not automatically removed until after the expiration period for an instant access virtual machine restore, which is 7 days by default. This behavior is intentional for the following reasons:

- To avoid data loss, since changes might have been made to the virtual machine during that time
- To provide you with the opportunity to fix the issue (for example, to free up space on the restore destination or choose a different datastore) and then take the appropriate action

When the cause of the failure is determined and/or fixed, you can use the **Instant Access Sessions** window of the UI to retry the migration, or save the data and delete the instant access virtual machine, as required. The section **Manage and monitor Instant Access Sessions** provides detailed information about these actions.

VMware knowledge base articles and product documentation

Additional VMware troubleshooting information is available at the VMware Knowledge Base and VMware Documentation websites.

Troubleshooting vSphere Plugin deployments

When investigating issues with the vSphere Plugin deployments, you might need to troubleshoot its deployment.

Troubleshoot vSphere Plugin deployments


In some circumstances, issues can occur during the deployment of the PowerProtect Data Manager **vSphere Plugin**.

About this task

If deployment of the **vSphere Plugin** fails, the plugin displays SSL errors or other errors such as `503 Service Not Available` or `No Healthy Upstream`, or you need to force the removal and re-installation of the plugin, perform the following steps:

Steps

1. In the PowerProtect Data Manager UI, go to **Infrastructure > Asset Sources**.
2. Select the vCenter asset source, and then click **Edit**.
3. Unselect **vSphere Plugin**, and then click **Save**.
4. Log in to the vCenter mob, for example, `http://vcenter.example.com/mob`.
5. Navigate to a new window to unregister the extension, for example, `http://vcenter.example.com/mob/?moid=ExtensionManager&method=unregisterExtension`
6. On this window, type `com.emc.dpsg.ppdm.plugin`, and then click **Invoke Method**.
7. In the PowerProtect Data Manager UI, go to **Infrastructure > Asset Sources**, select the vCenter, and click **Edit**.
8. Select **vSphere Plugin**, and then click **Save**.
9. Log out of the vCenter Server, and then log back into the vCenter.

 **NOTE:** If this is a newer vCenter server, a blue bar displays with a Refresh button. Click **Refresh**.

VMware knowledge base articles and product documentation

Additional VMware troubleshooting information is available at the VMware Knowledge Base and VMware Documentation websites.

Dell PowerProtect Cloud Snapshot Manager

Seamless snapshot management across multiple clouds for backup and disaster recovery

Essentials

- **Simple:** Automated discovery and protection of public clouds VMs, databases and block storage volumes based on policies. Set and forget simplicity.
- **Powerful SaaS:** Data Protection with multi-tenancy capabilities. Nothing to install, zero infrastructure cost
- **Multicloud:** Seamless management across multiple clouds (AWS, Azure & GCP)
- **Global Visibility and Control:** Dashboards and global reports enable enterprises to gain visibility and control into their cloud environments
- **Scalable:** Enterprise-grade solution for public cloud infrastructure protection, no matter how extensive the "sprawl." Advanced recovery options across regions and accounts for DR.
- **Data Sources for AWS:** EC2, EBS volumes, RDS DBs, Aurora, Redshift, DynamoDB
- **Data Sources for Azure:** VMs with managed disks and Blob storage containers
- **Data Sources for Google Cloud:** FLR capabilities on VMs

Software-as-a-Service cloud protection

The explosive growth of public cloud computing is transforming enterprise IT infrastructure. More and more enterprises are adopting a multi-cloud strategy that requires a seamless cloud-native data protection strategy. Organizations find it difficult to manage workloads and the proliferation of snapshots with the native tools offered by most cloud providers. Managing different tools across clouds to protect workloads running in different cloud is costly and can become overwhelming. What's required is an automated, enterprise-grade solution that provides global visibility and control over protection of workloads in the cloud.

PowerProtect Cloud Snapshot Manager is a SaaS component of Dell PowerProtect Data Manager and requires no hardware or installation. It automates the protection of cloud-native workloads via tag-based policies and leverages the underlying snapshot technology of the cloud. Customers use a single tool to discover, orchestrate, and automate the protection of workloads across AWS, Azure and Google Cloud Platform (GCP). With Cloud Snapshot Manager, you are able to control costs and avoid snapshot sprawl.

Simple, trusted and powerful data protection

Cloud Snapshot Manager is designed for any size cloud infrastructure and scales as your organization and data grows. It delivers a single interface that gives businesses global visibility, reporting and control of data protection activities across all their cloud accounts. Automated policy assignment of resources is essential to achieve auto-scaling in the cloud with a peace of mind that your resources are protected.



Additional value is delivered by taking advantage of PowerProtect DD Virtual Edition. With industry leading deduplication and low-cost cloud object storage, you can retain data longer while reducing costs.

Protection across multiple cloud accounts and regions

Cloud Snapshot Manager is breaking cloud silos and provides a better way to protect all workloads running in multiple clouds and accounts. From a single console, customers can discover, orchestrate, and automate the protection of workloads across AWS, Azure, and GCP. In addition, Cloud Snapshot Manager can be configured with policies to automate the copy of snapshots from one region to another to enable recoveries in case a disaster strikes.

Powerful policy-based backups

- Discovery of existing snapshots for better control over snapshot sprawl
- Seamless policy-based creation and deletion of snapshots
- Tag based protection policy assignment
- Application consistency via pre/post scripts
- Copy snapshot data to PowerProtect DD Virtual Edition (DDVE) based on policies in order to reduce storage costs

Recovery capabilities

- File Level Restore is provided AWS, Azure, GCP
- One click restore of VM with all its configuration
- Group Restore of many VMs
- Granular blob level restore
- Copy snapshot to an alternate account for added security and DR (AWS)
- Protection across multiple cloud accounts and regions
- Support incremental snapshot capability offered by Azure

Simple & scalable

- Global reports and dashboards providing visibility into health of backups for all accounts and supported clouds
- Auto scale to protect thousands of resources
- Audit logging and multi-tenancy
- REST API integration
- Support for Federated Identification with SAML



[Learn More](#) about Cloud
Snapshot Manager




[Contact](#) a Dell Technologies Expert

PowerProtect Data Manager 19.10

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
Supported Internet Protocol versions.....	11
Prerequisites.....	11
Firewall and port considerations.....	12
Role-based security.....	13
Data-in-flight encryption.....	13
PowerProtect Data Manager new deployment overview.....	13
PowerProtect Data Manager existing deployment overview.....	14
Chapter 2: Enabling the File System Agent.....	15
About the File System agent.....	15
Application agent and File System agent coexistence.....	15
File System agent prerequisites.....	17
File System agent limitations.....	18
Block-based backups.....	19
Best practices for file system backups.....	19
Configure the file system parallel backup setting.....	19
Configure asset multi-streaming for file-based backups.....	20
Protect an asset with the File System agent.....	21
Protect an asset in a Microsoft Windows Server clustered environment with the File System agent.....	21
Installing and uninstalling the File System agent on AIX.....	22
Install the File System agent on AIX.....	22
Uninstall the File System agent on AIX.....	23
Installing and updating the File System agent on Linux.....	24
Install the File System agent on Linux.....	24
Update the File System agent on Linux.....	26
Uninstall the File System agent on Linux.....	27
Recommission the File System agent on Linux or AIX.....	28
Installing and updating the File System agent on Windows.....	28
Install the File System agent on Windows.....	28
Update the File System agent on Windows.....	30
Uninstall the File System agent on Windows.....	32
Recommission the File System agent on Windows.....	33
Manage the PowerProtect agent service.....	33
About the PowerProtect agent service.....	33
Start, stop, or obtain the status of the PowerProtect agent service.....	34
Recovering the PowerProtect agent service from a disaster.....	35
Chapter 3: Managing Storage, Assets, and Protection.....	37
Manage the File System agent.....	37
View application agent details.....	38
Add protection storage.....	39

View the storage unit password.....	40
Enable an asset source.....	40
Disable an asset source.....	41
Delete an asset source.....	41
Discover a file system host.....	42
Adding a protection policy for File System protection.....	42
Add a protection policy for centralized File System protection.....	43
Add a protection policy for self-service File System protection.....	46
Add a policy to exclude assets from data protection operations.....	50
Cancel a File System agent backup or restore job.....	50
Add a service-level agreement.....	51
Extended retention.....	54
Edit the retention period for backup copies.....	56
Delete backup copies.....	57
Retry a failed backup copy deletion.....	57
Export data for deleted backup copies.....	58
Remove backup copies from the PowerProtect Data Manager database.....	58
Exclusion filters.....	59
Add an exclusion filter.....	59
Guidelines for exclusion filters.....	60
Edit or delete an exclusion filter.....	61
Apply an exclusion filter to a protection policy.....	61
Remove an exclusion filter from a protection policy.....	62
Centralized restore of a file system asset.....	62
Centralized image-level restore of a file system asset.....	63
Centralized file-level restore of a file system asset.....	64
Enable the File System agent after hostname change.....	65
Chapter 4: Performing Self-Service Backups and Restores with the File System Agent.....	67
Performing self-service backups of file systems.....	67
Performing self-service restore of a file system host.....	68
Using the ddfsadmin utility for file systems.....	68
Self-service image-level restore of file systems.....	68
Self-service file-level restore of file systems.....	69
Chapter 5: Performing Disaster Recovery with the File System Agent in Windows.....	70
Disaster recovery limitations.....	70
Preparing for disaster recovery.....	70
Gathering key information.....	70
Critical volumes in disaster recovery.....	70
Discover the assets to back up.....	71
Create a disaster recovery protection policy.....	71
Synchronize all clocks.....	72
Manually run a disaster recovery policy.....	72
Performing system-state recovery.....	72
Perform a system-state recovery.....	72
Recovering the Active Directory.....	73
Configure the client to boot into Directory Services Restore Mode.....	73
Recover the Active Directory from disaster recovery data.....	74

Authoritative and nonauthoritative Microsoft restores.....	74
Performing bare-metal recovery.....	75
Bare-metal recovery requirements.....	75
About the WinPE image.....	76
Using a custom WinPE image.....	76
Perform a bare-metal recovery.....	78
Performing a bare-metal recovery of Windows clusters.....	80
Performing application restores after bare-metal recovery.....	81
Appendix A: File System Best Practices and Troubleshooting.....	82
Installation and operation.....	82
Backups.....	83
Disaster recovery.....	85
Restores.....	86
Storage units.....	87

Preface

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. In many cases the user interface has not yet been updated to reflect this change.

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.

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 employee will update the document as necessary.

Purpose

This document describes how to configure and administer the Dell EMC PowerProtect Data Manager software 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 software.

Revision history

The following table presents the revision history of this document.

Table 1. Revision history

Revision	Date	Description
02	April 4, 2022	This revision includes the following updates: <ul style="list-style-type: none">Added information about protecting assets in a clustered environment.

Table 1. Revision history (continued)

Revision	Date	Description
		<ul style="list-style-type: none"> Updated the filename for the Window File System agent installer.
01	March 22, 2022	Initial release of this document for PowerProtect Data Manager version 19.10.

Compatibility information

Software compatibility information for the PowerProtect Data Manager software is provided at 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 Administration and User Guide</i>	Describes how to configure the software.
<i>PowerProtect Data Manager Deployment Guide</i>	Describes how to deploy the software.
<i>PowerProtect Data Manager Licensing Guide</i>	Describes how to license the software.
<i>PowerProtect Data Manager Release Notes</i>	Contains information on 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 Dell EMC 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.
<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.
<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.

Table 2. Related documentation (continued)

Title	Content
<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.
<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.
<i>PowerProtect Data Manager Disaster Recovery Best Practices Guide</i>	Provides guidance and best practices for a PowerProtect Data Manager server disaster-recovery solution.
PowerProtect Data Manager Public REST API documentation	Contains the PowerProtect Data Manager 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 Dell EMC 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

- The Customer Support website

- The Community Network

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.

Knowledgebase

The Knowledgebase contains applicable solutions that you can search for either by solution number (for example, KB000xxxxxx) or by keyword.

To search the Knowledgebase:

1. Go to the Customer Support website.
2. On the **Support** tab, click **Knowledge Base**.
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 DPAD.Doc.Feedback@ems.com.

PowerProtect Data Manager File System Agent Overview

Topics:

- PowerProtect Data Manager overview
- Supported Internet Protocol versions
- Prerequisites
- Firewall and port considerations
- Role-based security
- Data-in-flight encryption
- PowerProtect Data Manager new deployment overview
- PowerProtect Data Manager existing deployment overview

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.

Supported Internet Protocol versions

PowerProtect Data Manager only supports the use of IPv4 addresses.

Using an IPv6 address can result in errors or other unexpected behavior. When configuring devices to connect over the network with PowerProtect Data Manager, use only IPv4 addresses.

Prerequisites

Ensure that your environment meets the requirements for a new deployment or update of PowerProtect Data Manager.

Requirements:

- ① **NOTE:** The most up-to-date software compatibility information for the PowerProtect Data Manager software and the application agents is provided in 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.
 - ① **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 19.10 or earlier is required.
- License: A trial license is provided with the PowerProtect Data Manager software. Dell EMC Data Protection Suite Applications, Backup, and Enterprise customers can contact Dell EMC Licensing Support for assistance with a permanent PowerProtect Data Manager license.

- Large environments require multiple PowerProtect Data Manager instances. Contact Champions.eCDM@emc.com for assistance with sizing requests.
- The PowerProtect Data Manager 19.10 download file requires the following:
 - ESXi version 6.5, 6.7, or 7.0.
 - Eight vCPUs, 18 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 7000 is open between PowerProtect Data Manager and the application agent hosts.
- VMware ESXi server that hosts PowerProtect Data Manager meets the following minimum system requirements:
 - 10 CPU cores
 - 18 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

Firewall and port considerations

The latest version of the *PowerProtect Data Manager Security Configuration Guide* provides more details about the port requirements.

Table 4. PowerProtect Data Manager port requirements

Description	Communication	Port
SSH communications	Bi-directional communication between the SSH client and the PowerProtect Data Manager appliance.	22 TCP/UDP
Microsoft SQL Server, Oracle, Microsoft Exchange Server, SAP HANA, File System	Bi-directional communication between the PowerProtect Data Manager agent and the PowerProtect Data Manager appliance. Requirement applies to Application Direct and VM Direct.	7000 TCP
REST Server	Bi-directional communication between the HTTP client and the PowerProtect Data Manager appliance.	8443 TCP
RESTAPI Server - VM Direct	Bi-directional communication between the PowerProtect Data Manager agent and the PowerProtect Data Manager appliance. Requirement applies to Microsoft SQL Server VM application-aware.	8443 TCP
UI redirect	Inbound only.	80 TCP 443
LDAP	Outbound only.	389 TCP/UDP 636 TCP
Discovery (devices)	Outbound between the PowerProtect Data Manager appliance and the device.	3009 TCP—Storage Direct and DD system 5989 TCP—SMI-S 443 TCP—XtremIO 7225 TCP—RecoverPoint

Table 4. PowerProtect Data Manager port requirements (continued)

Description	Communication	Port
PowerProtect Data Manager agent	Bi-directional communication between the database hosts and the PowerProtect Data Manager appliance. This requirement applies to both Application Direct and VM Direct.	7000 TCP
Embedded VM Direct service	Outbound.	9090 TCP

Role-based security

PowerProtect Data Manager provides predefined user roles that control access to areas of the user interface and to protected operations. Some of the functionality in this guide 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.

Data-in-flight encryption

PowerProtect Data Manager provides centralized management of backup and restore encryption for application agents. Backup and restore encryption is supported for both centralized and self-service operations where applicable.

You can ensure that the backup and restore content is encrypted when read on the source system, transmitted in encrypted form, and then decrypted before it is saved on the destination storage. This prevents another party from intercepting private data.

PowerProtect Data Manager only supports encryption in-flight for Microsoft SQL Server with Application Direct, Microsoft Exchange Server, Oracle RMAN, File System, SAP HANA, and network attached storage (NAS) workloads. This is a global setting that is applicable to all supported workloads. 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 Administration and User 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 install 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. Install 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 install PowerProtect Data Manager with the File System agent in 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. Install PowerProtect Data Manager from the downloaded 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.

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 PowerProtect agent service

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:

- 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	<ol style="list-style-type: none"> 1. New installation of both agents is supported with: <ol style="list-style-type: none"> a. Microsoft SQL Server agent, Microsoft Exchange Server agent, Oracle RMAN agent, or SAP HANA agent in stand-alone or managed mode. b. File System agent in managed mode. 2. New installation of an agent is supported in managed mode with an existing agent in stand-alone mode. 3. New installation of an agent is supported in stand-alone mode with an existing agent in managed mode. 4. Repair of an existing agent installation is supported. 5. 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	<ol style="list-style-type: none"> 1. Registration of an installed agent to the PowerProtect Data Manager server is supported. 2. 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	<ol style="list-style-type: none"> 1. Protection policy creation is supported on all registered agents. 2. All scheduled protection policy backups are supported on both agents as per individual protection policies. 3. Self-service backups are supported on both agents. 4. Compliance is supported on both agents as per the individual Service Level Agreements (SLAs). 5. 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 the PowerProtect Data Manager network can see and resolve each other. If PowerProtect Data Manager and the File System agent are registered in different domains, you must add IP address and FQDN entries on both the client and server.
 1. Browse to the path of the hosts file. For example, on Windows `C:\Windows\System32\drivers\etc\hosts`, and on Linux `/etc/hosts`.
 2. Add an entry to the hosts file, as in the following:

```
IP address      FQDN            common name
10.10.100.100  yourdomain.com yourdomain
```

- LVM/VxVM partitions or volumes are supported.
- Physical 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.
- 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/narbbb.ko: Required key not available` appears. As a workaround, you can disable the Secure Boot option.
- Ensure that the file system has the `/etc/fstab` entry. Without the `/etc/fstab` entry, 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 the `cat /proc/sys/kernel/kptr_restrict` file to verify that permissions to install the driver are set. If the value is set to `2`, then there is a restriction 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.

See the Operating System documentation for the installation procedure.

Additional AIX prerequisites

- File system discovery requires a JFS or JFS2 file system type.
- Ensure that IBM XL C/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 and Linux limitations

- File System agent block-based backups exclude the following:
 - Application files such as Microsoft SQL Server and Microsoft Exchange Server files.
 - ① **NOTE:** For file-based backups, application data such as Microsoft SQL Server, Hyper-V, and Microsoft Exchange Server data is backed up.
 - HyperVisor files. The File System agent is installed primarily in the guest operating system for the backup of guest file system volumes, and is not dependent on the underlying HyperVisor.
 - 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.
- The File System agent supports operating systems in only the following languages:
 - English — full support.
 - Japanese — the File System agent can be used for all protection and restore operations on a Japanese-language operating system, provided that volume asset names are in English. Only file and folder names can be in Japanese.
- If a Windows or Linux file system host is unregistered from PowerProtect Data Manager and then re-registered 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.
- If the Bytes of sector sizes of the source and target volumes are different, PowerProtect Data Manager does not support block-based image recoveries. For example, you cannot perform a block-based image recovery of a volume that has 4096 as the Bytes of sector size to a volume that has 512 as the Bytes of sector size, 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.

Linux limitations

- 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/nsrbbb.ko: Required key not available appears. As a workaround, you can disable the Secure Boot option.

- A file system backup might fail with the error `Insufficient space exists in the volume group for creating shadow of the volume` when there is not enough space in the volume group for a block-based backup to succeed. Each volume group on LVM2 or VxVM must have at least 10% free space.
- On Linux, performing an image-level restore of a block-based backup volume to an alternate location changes the GUID for the volume. As a result, PowerProtect Data Manager displays duplicate assets on the **Infrastructure > Assets** pane, where the older asset has a status of **Deleted/Not Detected**. To ensure continued protection, replace the old asset in the protection policy with the new asset.

AIX limitations

- On AIX, PowerProtect Data Manager supports only file-based backups. Block-based backups are not supported.

Block-based backups

Block-based backups provide instant access to the backups. The block-based backups enable you to mount the backups by using the same file systems that you used to back up the data.

The File System agent's block-based backups support the following capabilities:

- Mounting of a backup as a file system
- Mounting of an incremental backup
- Sparse backup support
- Backups of operating system-deduplicated file systems as source volumes on Windows
- Forever virtual full backups to DD
- DD retention lock
- Recoveries from DD

Block-based backups are useful for datasets that are under 10 TB with a single volume under 5 TB, and a daily change rate under 5%.

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:

- $\text{Data backup speed} = \text{asset size} / \text{time taken to back up the asset}$
- $\text{Total backup size} = \text{sum of the total size of the assets to be backed up for a full backup}$
- $\text{Maximum backup window} = \text{total backup size} / \text{data backup speed}$
- $\text{Optimal backup window} = \text{maximum backup window} / \text{parallelism value}$
- 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.

Configure the file system parallel backup setting

PowerProtect Data Manager enables you to run file system backups in parallel to reduce the time taken for backups. This setting defines the maximum concurrent network sessions from the client to the DD system at any given time. You can specify the number of streams to use for the backup in the configuration file `./ddfssv.fsagent.config` or through the self-service

CLI. However, it is recommended that you set the parallelism value in the configuration file as the parallelism value provided in the configuration file takes precedence over the parallelism value that is provided in the CLI.

NOTE: Backup parallelism is only available on supported Windows systems. Since the parallelism setting is defined at the host level, you must set the parallelism setting on every Windows host where parallel file system backups are enabled.

To specify the number of streams to use for the backup, you can set the file system parameter `--max-host-streams` in the `.ddfssv.fsagentconfig` file that is located in the `C:\Program Files\DPSAPPS\fsagent\settings` directory on the file system host. This value must be an integer. The default value is 8.

For example, if you set the `--max-host-streams` parameter value to 6, the File System agent backups run with 6 streams:

```
--max-host-streams=6
```

You can also use the command-line option `-M` to specify the number of streams to use for the backup. For example:

```
ddfssv -LL -l INCR -z -M 4 -a DFA_SI=TRUE  
-a DFA_SI_USE_DD=TRUE -a DFA_SI_DD_HOST=protection_storage_system_ip_address  
-a DFA_SI_DD_USER=protection_storage_user -a  
DFA_SI_DD_PASSWORD=protection_storage_system_password -a  
DFA_SI_DEVICE_PATH=protection_storage_device_path I:
```

where:

- a "DFA_SI_DD_HOST=protection_storage_system_ip_address"
Specifies the name of the DD server that contains the storage unit where you want to back up the databases.
- a "DFA_SI_DD_PASSWORD=protection_storage_system_password"
Specifies the password of the protection storage system user.
- a "DFA_SI_DEVICE_PATH=protection_storage_device_path"
Specifies the name and the path of the protection storage device where you want to direct the backup.
- a "DFA_SI_DD_USER=protection_storage_user"
Specifies the protection storage username. For example, sysadmin.

Recommendations for optimal performance:

- Tune the parallelism value based on the resources that are available on your system, including the CPU, the number of assets, and the connection bandwidth to the DD system.
- Set the parallelism value equal to the number of CPU cores on the host.
- Run backups in parallel for several assets that are small in size, as opposed to running backups in parallel for a few assets that are large in size.

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, `.ddfssv.fsagentconfig` under the `Settings` folder.

Multi-streaming may not yield desired results in the following scenarios:

- System has limited CPUs/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 [Add protection storage](#).
 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 [Adding a protection policy for File System protection](#).
- 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 (CSVs) 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 [Add protection storage](#).
 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:** Standalone 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 Adding a protection policy for File System protection.

NOTE: The backup of a cluster asset is routed through the node on which the asset or volume is active.

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.
- 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:
 - a. Select **Agent Downloads** from the **System Settings** menu.
 - b. Select the File System agent download package for AIX, `fsagent1910_aixpower72.tar.gz`.
 - c. 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:
 - a. `gunzip fsagent1910_aixpower72.tar.gz`
 - b. `tar -xvf fsagent1910_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-agentsvc.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. To enable the PowerProtect Data Manager communication 7000 TCP port, run the `/opt/dpsapps/agentsvc/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.

NOTE:

The `--server` option is used to specify the PowerProtect server IP for registration, and is mandatory for silent install.

Use the following commands to perform a silent installation on AIX:

- For silent installation, including registration of the agent with the PowerProtect server, type: `install.sh --silent-install --server=<PowerProtect_server_IP>`
- To run the installer in debug mode during silent installation, type: `install.sh --debug --silent-install --server=<PowerProtect_server_IP>`
- To skip installation of the block-based backup driver, type: `install.sh --skip-driver --silent-install --server=<PowerProtect_server_IP>`
- To enable the PowerProtect Data Manager communication 7000 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 `fsagent1910_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**.

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-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.
- 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. Select **Agent Downloads** from the **System Settings** menu.
 - b. Select the File System agent download package for Linux, `fsagent1910_linux_x86_64.tar.gz`.
 - c. 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 Linux is not supported.
2. Untar the installer by running `tar -xvf fsagent1910_linux_x86_64.tar.gz`.
Then run the command `cd fsagent` to change the current working directory to the extracted path.
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`.
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-agent.svc.rpm` or `powerprotect-agent.svc.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.
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. To enable the PowerProtect Data Manager communication 7000 TCP port, run the `/opt/dpsapps/agentsvc/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/nsrbbb.ko: Permission denied.`
3. To check the audit log, run: `ausearch -c '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/nsrbbb-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 -c '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.

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 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 server, type: `install.sh --silent-install --server=<PowerProtect_server_IP>`
- To run the installer in debug mode during silent installation, type: `install.sh --debug --silent-install --server=<PowerProtect_server_IP>`
- To skip installation of the block-based backup driver, type: `install.sh --skip-driver --silent-install --server=<PowerProtect_server_IP>`
- To enable the PowerProtect Data Manager communication 7000 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 interactive or 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.

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.

Update the File System agent on Linux in interactive mode

Use this interactive procedure to update 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.

Steps

1. In the PowerProtect Data Manager UI:
 - a. Select **Agent Downloads** from the **System Settings** menu.
 - b. Select the File System agent download package for Linux, `fsagent1910_linux_x86_64.tar.gz`.
 - c. 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 Linux is not supported.

2. Untar the installer by running `tar xvf fsagent1910_linux_x86_64.tar.gz`.

Then run the command `cd fsagent` to change the current working directory to the extracted path.

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 `install.sh` script.

The following `.rpm` or `.deb` files are installed as part of the script:

- `powerprotect-agent.svc.rpm` or `powerprotect-agent.svc.deb`—Installs or updates the agent service component for 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.

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.

5. To enable the PowerProtect Data Manager communication 7000 TCP port, 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.

NOTE: If the firewall rules are not applied, restart the firewall.

Update the File System agent on Linux in silent mode

Use the following commands to perform a silent update of the File System agent on Linux:

NOTE: To verify the authenticity and integrity of the RPM files prior to the update step, follow the instructions in the *PowerProtect Data Manager Security Configuration Guide*.

- To run the silent agent update, type: `install.sh --force-upgrade`

NOTE: The `--force-upgrade` command line option does not prompt users to update other agents so that they are running the same version as the File System agent. Ensure that other agents are running the same version as the File System agent, otherwise those agents will not work.

- To run the silent update and register the agent with a new PowerProtect server, type: `install.sh --force-upgrade --server=PowerProtect Server IP`
- For silent update in debugging mode, type: `install.sh --force-upgrade --debug`

NOTE: During silent installation/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/agent.svc/configfw.sh` script as the root user.

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 `fsagent.1910.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-agent.svc. Do you wish to
uninstall powerprotect-agent.svc? (y/n)
```

3. Type **Y** to confirm that you want to uninstall `powerprotect-agent.svc`.

NOTE:

If you type **N**, the .rte files (ppdm_fsagent.rte, ppdm-bbbvt.rte and ppdm-fsagent.rte) are uninstalled. However, powerprotect-agentSvc remains in an installed state.
There is no silent uninstall procedure on Linux.

Recommission the File System agent on Linux or AIX

You can use the procedure in this topic to onboard the decommissioned File System agent to the same PowerProtect Data Manager server.

NOTE:

You can run the `install.sh` or `register.sh` script to register and recommission 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 recommission 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 recommissioned.

To use the `register.sh` script to register and recommission 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. Select **Agent Downloads** from the **System Settings** menu.
 - b. Select the File System agent download package for Windows, `fsagent1910_win_x64.zip`.
 - c. Download the package in the location that you want to install the File System agent.
2. Run the `fsagent-19.10.0.0.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 7000, ensure that the **Configure the Windows Firewall** option is selected under **Common Core Components**. This option is selected by default.

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 port 7000 on the File System agent host and port 8443 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.

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 re-register 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-19.10.0.0.exe /s PPDMSHostName=<PPDM_server_IP>
```

To perform the silent installation to a different path, run:

```
fsagent-19.10.0.0.exe /s PPDMSHostName=<PPDM_server_IP>  
ProductInstallPath="D:\alternate_path"
```

To enable the PowerProtect Data Manager communications port 7000, if Windows firewall rules have not been previously configured, run:

```
fsagent-19.10.0.0.exe /s PPDMSHostName=<PPDM_server_IP> EnableFirewallRules=1
```

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 port 7000 on the File System agent host and port 8443 on PowerProtect Data Manager to be open bidirectionally. These ports enable communication between the File System agent and PowerProtect Data Manager.

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: `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, 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.

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.
- Downloaded the File System agent software package.

Steps

1. In the PowerProtect Data Manager UI:
 - a. Select **Agent Downloads** from the **System Settings** menu.
 - b. Select the File System agent download package for Windows, `fsagent1910_win_x64.zip`.
 - c. Download the package in the location that you want to install the File System agent.
2. Run the `fsagent-19.10.0.0.exe` program.
3. Follow the update steps in the wizard to provide the installation location and the PowerProtect Data Manager server IP address.

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.

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 7000 is not enabled, you can change the firewall rule setting now as part of this update by selecting the **Configure the Windows Firewall** option under **Common Core Components**.

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 port 7000 on the File System agent host and port 8443 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.

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.
 - `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 FQDN of the client, the installation completes successfully but the registration fails. To re-register 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-19.10.0.0.exe /s PPDMHostName=<PPDM_server_IP> ForceUpgrade=1
```


To perform the silent update to a different path, run:

```
fsagent-19.10.0.0.exe /s PPDMServerName=<PPDM_server_IP> ForceUpgrade=1  
ProductInstallPath="D:\alternate_path"
```

To enable the PowerProtect Data Manager communications port 7000, if Windows firewall rules have not been previously configured, run:

```
fsagent-19.10.0.0.exe /s PPDMServerName=<PPDM_server_IP> EnableFirewallRules=1  
ForceUpgrade=1
```

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 port 7000 on the File System agent host and port 8443 on PowerProtect Data Manager to be open bidirectionally. These ports enable communication between the File System agent and PowerProtect Data Manager.

The change in `EnableFirewallRules` configuration will take place only on first-time installation or upgrade of the agent service.

NOTE: `PPDMServerName` 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 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-19.10.0.0.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 PowerProtect Data Manager communications port 7000 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-19.10.0.0.exe /s /uninstall
```

- To perform a silent uninstall while also uninstalling common components, run:

```
fsagent-19.10.0.0.exe /s /uninstall UnInstallPPCMAgent="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.

Recommission the File System agent on Windows

You can use the procedure in this topic to onboard the decommissioned File System agent to the same PowerProtect Data Manager server.

NOTE:

You can use the **Modify** option under **Add/Remove programs** for the File System agent to recommission 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 Windows.

Alternatively, you can use the `register.bat` script to register and recommission the File System agent with PowerProtect Data Manager, as follows:

```
<Install_folder>\AgentService\register.bat --enable
```

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 standalone 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.
 - ① **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.
 - ① **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.

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.

- ① **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 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 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 clearing 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.

Managing Storage, Assets, and Protection

Topics:

- Manage the File System agent
- Add protection storage
- Enable an asset source
- Discover a file system host
- Adding a protection policy for File System protection
- Cancel a File System agent backup or restore job
- Add a service-level agreement
- Extended retention
- Edit the retention period for backup copies
- Delete backup copies
- Exclusion filters
- Centralized restore of a file system asset
- Enable the File System agent after hostname change

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 Windows or Linux 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. In the **Add Application/FS Agent** window, select one of the following options:

- **Add FQDN**

Perform the following steps:

- a. Type the fully qualified domain name (FQDN) 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 the **Choose File** icon.

NOTE: The contents of the .csv file must be in the following format, for example:

```
"ppdm.dell.com"
```

```
"ppdm2.emc.com"  
"ppdm.dell EMC.com"
```

The **Explorer** window appears.

- b. Select the `.csv` file, and then click **Open**.

The file appears in the **Application/FS Agents** 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**

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. You can also view information about scheduled updates by clicking the arrow at the end of the row. 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	The name of the application agent host.
IP	The IP address of the application agent host.
Registration Status	The registration status of the application agent: <ul style="list-style-type: none">• Waiting• Pending• Registered• Approved• Rejected• Expired• Acceptingocerts
OS	The operating system of the application agent host.
Agent Types	The application agent type.

Table 5. Application agent information (continued)

Column	Description
Current Version	The current version of the application agent.
Update Status	<p>The update status of the application agent host:</p> <ul style="list-style-type: none"> • Up to Date • Available • Scheduled • In Progress • Failed • Not Supported <p>NOTE: In this release of PowerProtect Data Manager, the update status of application agents is always Up to Date.</p>
Registered Date	The date that the application agent was registered in PowerProtect Data Manager.
Created Date	The creation date of the application agent.

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.

Add protection storage

Add and configure a storage system to use as a target for protection policies. Adding a storage system requires the admin role.

About this task

The *PowerProtect Data Manager Administration and User Guide* provides more information about protection storage and related concepts:

- High availability options
- Smart Scale system pools, a single interface to a flexible group of pool members
- Working with protection storage
- Working with storage units

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. 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.
5. 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**.

- 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**.
- 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.
- 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**.
- 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.
 - In the **Storage** window, select the storage system from the table.
 - Click **More Actions > Set Location**.
The **Set Location** window appears.
 - Click **Add** in the **Location** list.
The **Add Location** window appears.
 - In the **Name** field, type a location name for the asset, and click **Save**.

Results

PowerProtect Data Manager displays external DD systems only in the **Storage** window **Name** column; PowerProtect Data Manager displays DD Management Center storage types in the **Managed By** column.

The *PowerProtect Data Manager Administration and User Guide* provides more information about working with storage units, such as the relationships between storage units and policies, and applicable limitations.

View the storage unit password

PowerProtect Data Manager provides a script to retrieve the password for a storage unit that you configured as a backup target.

Prerequisites

This task requires the name of the PowerProtect DD MTree where the storage unit resides.

Steps

- Connect to the PowerProtect Data Manager console as an admin user.
- Navigate to the `/usr/local/brs/puppet/scripts` directory.
- Obtain the storage unit password by typing the following command:

```
python get_dd_mtree_credential.py MTree-name
```

For example:

```
python get_dd_mtree_credential.py ppdm-1910-blrv034d018-75914
-----PowerProtect DD MTree credential-----
Full MTree path: /data/coll/ppdm-1910-blrv034d018-75914
User name: ppdm-1910-blrv034d018-75914
Password: lWXT#DC93m={XV+K
-----
```

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 sources 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.

- ① **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

- ① **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 virtual machines, then you cannot disable the vCenter 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 can be deleted automatically or manually. The *PowerProtect Data Manager Administration and User 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. Note that each time you initiate a discovery process, the resources that are discovered and those that are handling the discovery impact system performance.

4. Click **Yes**.

Results

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 the  icons 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.

Adding a protection policy for File System protection

Use the PowerProtect Data Manager UI to add a protection policy to protect File System data.

Review the prerequisites in the section File System agent prerequisites.

Before you perform a backup on a weekly or monthly schedule from the protection policy, set the PowerProtect Data Manager time zone to the local time zone. If the PowerProtect Data Manager time zone is not set to the local time zone, the weekly or monthly backup still runs but triggers based on the PowerProtect Data Manager time zone.

The *PowerProtect Data Manager Administration and User Guide* provides more information about working with storage units, such as the relationships between storage units and policies, and applicable limitations.

If applicable, complete all of the virtual network configuration tasks before you assign any virtual networks to the protection policy. The *PowerProtect Data Manager Administration and User Guide* provides more information.

Policy types and purposes

Each File System protection policy has a type that defines the purpose of the policy. When you add a protection policy, select one of the available types:

- **Centralized Protection**—Use PowerProtect Data Manager to manage all aspects of protection.
- **Self-Service Protection**—Use the File System agent to create local backup protection. PowerProtect Data Manager creates a protection policy and manages extra stages.
- **Exclusion**—If there are assets within the protection policy that you plan to exclude from data protection operations.


For clarity, this guide provides separate instructions to add a protection policy of each type, even though some of the steps overlap. The terminology for each task differs slightly where the available stages differ between policy types. Complete only the task that applies to your goal.

Add a protection policy for centralized File System protection


With centralized protection, PowerProtect Data Manager manages all aspects of the protection process. The process of adding a protection policy is similar for all policy types. However, these instructions contain only elements and options that appear when you select the centralized protection policy type.



Prerequisites

Review the prerequisites in *Adding a protection policy for File System protection*.

-  **NOTE:** To enable replication, ensure that you add a remote DD system as the replication location. The *PowerProtect Data Manager Administration and User Guide* provides detailed instructions about adding a remote DD system.

Steps

1. Select **Protection > Protection Policies**.
The **Protection Policy** window appears.
2. Click **Add**.
The **Add Policy** window appears.
3. In the **Type** page, specify the new protection policies group fields. For example, if you are creating a protection policy for daily backups in the Windows 2012 Server:
 - a. In the **Name** field, specify the name of the protection policy. For example, **File System Prod**.
 **NOTE:** The name that you specify here becomes part of the DD MTree entry.
 - b. In the **Description** field, specify a short description of the protection policy. For example, **File System Prod Daily Backups**.
 - c. In the **Type** field, select **File System**.
 - d. Click **Next**.
The **Purpose** page appears.
4. To manage all protection centrally, click **Centralized Protection**.
5. Click **Next**.
The **Assets** page appears.
6. 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 the   icons 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.
7. Click **Next**.
The **File Exclusions** page appears.
8. Optionally, to enable exclusions, click **Enable**.
 - a. Select one or more filters to apply, provide the parameters, and click **Add Filter**.
Click **Add a saved filter** to use an existing filter or group of filters as a template.

i | **NOTE:** Add an exclusion filter provides more details about exclusion filters.

- b. Enter a name and description for the filter and click **Save**.

The **Objectives** page appears.

9. On the **Objectives** page, 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 Administration and User Guide* provides instructions.

10. Click **Add** under **Primary Backup**.

The **Add Primary Backup** dialog appears.

11. On the **Schedules** pane of the **Add Primary Backup** dialog:

- a. Specify the following fields to schedule the synthetic full backup of this protection policy:

- **Create a Synthetic Full...**—Specify how often to create a synthetic full backup. A synthetic full backup copies only the changed data since the last backup to create a new full backup.

- **Retain For**—Specify the retention period for the synthetic full backup.

You can extend the retention period for the latest primary backup copy by using the **Extend Retention** schedule. For example, your regular schedule for daily backups can use a retention period of 30 days. However, you can apply extended retention to keep the full backups taken on Mondays for 10 weeks. Extended retention provides more information about extended retention.

- **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.

i | **NOTE:** Any backups started before the **End Time** occurs continue until completion.

- Click **Save** to save and collapse the backup schedule.

- b. Click **Add Backup** if you want to periodically force a full (level 0) backup, and then specify the following fields to schedule the full backup of this protection policy:

i | **NOTE:** When you select this option, the backup chain is reset.

- **Create a Full...**—Specify whether you want to create a weekly or monthly full backup.

- **Repeat on**—Depending on the frequency of the full backup schedule, specify the day of the week or the date of the month that the full backup occurs.

- **Retain For**—

△ 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.

- **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.

i | **NOTE:** Any backups started before the **End Time** occurs continue until completion.

- Click **Save** to save and collapse the backup schedule.

12. On the **Target** pane of the **Add Primary Backup** dialog, specify the following fields:

- a. **Storage Name**—Select a backup destination from the list of existing DD systems, or select **Add** to add a system and complete the details in the **Storage Target** window.

i | **NOTE:** The **Space** field indicates the total amount of space, and the percentage of available space, on the storage system.

- b. **Storage Unit**—Select whether this protection policy should use a **New** storage unit on the selected DD 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-ppdm-daily-123ab (300 GB/1 TB)**

When you select **New**, a new storage unit in the format *policy name host name unique identifier* is created in the storage system upon policy completion. For example, **testvmpolicy-ppdm-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 on the selected system. PowerProtect Data Manager uses Governance mode for retention locking, which means that the lock can be reverted at any time if necessary. Toggling the **Retention Lock** slider on or off applies to the current backup copy only, and does not impact the retention lock setting for existing backup copies.

NOTE: If you enable **Retention Lock** for a replicated backup, set the **Retain For** field in the **Add Replication** backup schedule dialog to a period that is at least as long as the retention lock period for the primary backup, so that the replicated backup does not expire before the primary backup.

- e. **SLA**—Select an existing service level agreement that you want to apply to this schedule from the list, or select **Add** to create an SLA within the **Add Service Level Agreement** wizard.

The *PowerProtect Data Manager Administration and User Guide* provides instructions.

13. Click **Save** to save your changes and return to the **Objectives** page.

The **Objectives** page updates to display the name and location of the target storage system under **Primary Backup**.

NOTE: After completing a backup schedule, you can change any schedule details by clicking **Edit** next to the schedule.

14. Optionally, extend the retention period for the latest primary backup copy:

Extended retention provides more information about **Extend Retention** functionality.

- a. Click **Extend Retention** next to **Primary Backup**. An entry for **Extend Retention** is created below **Primary Backup**.

- b. Under **Extend Retention**, click **Add**. The **Add Extended Retention** dialog appears.

- c. **Retain the next scheduled full copy every...**—Specify a weekly, monthly, or yearly recurrence for the extended retention backup schedule.

- d. **Repeat on**—Depending on the frequency of the full backup schedule, specify the day of the week, the date of the month, or the date of the year that the extended retention backup occurs.

- e. **Retain For**—Specify the retention period for the backup. You can retain an extended retention backup for a maximum of 70 years.

- f. Click **Save** to save your changes and return to the **Objectives** page.

15. Optionally, replicate these backups to a remote storage system:

- a. Click **Replicate** next to **Primary Backup** or **Extend Retention**. An entry for **Replicate** is created to the right of the primary or extended retention backup.

NOTE:

PowerProtect Data Manager supports replicating an extended retention backup only if the primary backup already has one or more replication stages. Also, for replication of an extended retention backup, you can only select from the DDs to which the primary stage replicates.

For example, if there are 6 DDs available (DD001-DD006), and the primary backup is on DD0001:

- Replicate1, which is based on the primary backup, replicates to DD002.
- Replicate2, which is based on the primary backup, replicates to DD003.
- Extended retention backup is backed up to DD001.
- Replicate3, which is based on the extended retention backup, must replicate to DD002 or DD003.

- b. Under **Replicate**, click **Add**. The **Add Replication** dialog appears.

- c. Complete the schedule details in the **Add Replication** dialog.

NOTE: The schedule frequency can be every day, week, month, or x hours for replication of the primary backup, and every day, week, month, year, or x hours for replication of the extended retention backup. 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 1 to 12 hours.

All replication copies of the primary backup schedule use the same retention period, and by default, this retention period is inherited from the **Retain For** value of the synthetic full backup schedule. To specify a different retention period for all of the replication copies of this primary backup schedule, click **Edit**, change the value in the **Retain For** field, and then click **Save**. This retention period will be applied to all of the replicated copies (synthetic full and full) of this primary backup schedule.

16. Optionally, add a cloud stage for a primary, replication, or extended retention schedule to move backups from DD storage to the cloud tier:

NOTE: To move a backup or replica to the cloud tier, schedules must have a retention time of 14 days or more. PowerProtect Data Manager also requires the discovery of a DD system with a configured cloud unit.

- a. Click **Cloud Tier** next to **Primary Backup** or **Extend Retention** or, if adding a cloud stage for a replication schedule that you have added, click **Cloud Tier** under **Replicate**. An entry for **Cloud Tier** is created to the right of the primary or extend retention backup schedule, or below the replication schedule.
- b. Under the entry for **Cloud Tier**, click **Add**.
The **Add Cloud Tier Backup** dialog appears, with summary schedule information for the parent node to indicate whether you are adding this Cloud Tier stage for the primary backup schedule, the extended retention backup schedule, or the replication schedule.
- c. Complete the schedule details in the **Add Cloud Tier Backup** dialog, and then click **Save** to save your changes and return to the **Objectives** page.
The *PowerProtect Data Manager Administration and User Guide* provides detailed instructions for adding a cloud stage for a primary, replication, or extended retention schedule.

17. Click **Next**.

The **Options** page appears.

18. On the **Options** page, select the additional option if required for the policy:

- **Troubleshooting**—Select this option to enable the debug logs for troubleshooting at higher debug levels. 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: Overriding the bug level in this manner may result in larger logs that may slow backup operations.

In Windows environments, the impacted logs include:

- o `FSAgentInstallPath\logs\vsscr.log`
- o `FSAgentInstallPath\logs\nsriscsi.log`
- o `FSAgentInstallPath\logs\nsriscsi_***.log`
- o `FSAgentInstallPath\logs\nsrwriter.log`
- o `FSAgentInstallPath\logs\ddfscon.***.log`
- o `FSAgentInstallPath\logs\ddfscon_***.log`
- o `FSAgentInstallPath\logs\ddfsav.log`
- o `FSAgentInstallPath\logs\ddfsav_***.log`
- o `FSAgentInstallPath\logs\ddfsrc_***.log`

In Linux environments, the impacted logs include:

- o `/opt/dpsapps/fsagent/logs/nriscsi.log`
- o `/opt/dpsapps/fsagent/logs/ddfacon.***.log`
- o `/opt/dpsapps/fsagent/logs/ddfasv.log`

NOTE: 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.

19. Click **Next**.

The **Summary** page appears.

20. 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.

When the new protection policy is created and assets are added to the protection policy, PowerProtect Data Manager performs backups according to the backup schedule.

21. 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.

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 PowerProtect DD.

Add a protection policy for self-service File System protection

With self-service protection, the agent or host controls the primary backup process and PowerProtect Data Manager manages other aspects of the protection process. The process of adding a protection policy is similar for all policy types. However, these instructions contain only elements and options that appear when you select the self-service protection policy type.



Prerequisites

Review the prerequisites in *Adding a protection policy for File System protection*.

- ① **NOTE:** To enable replication, ensure that you add a remote DD system as the replication location. The *PowerProtect Data Manager Administration and User Guide* provides detailed instructions about adding a remote DD system.

Steps

1. Select **Protection > Protection Policies**.
The **Protection Policy** window appears.
2. Click **Add**.
The **Add Policy** window appears.
3. In the **Type** page, specify the new protection policies group fields. For example, if you are creating a protection policy for daily backups in the Windows 2012 Server:
 - a. In the **Name** field, specify the name of the protection policy. For example, **File System Prod**.
① **NOTE:** The name that you specify here becomes part of the DD MTree entry.
 - b. In the **Description** field, specify a short description of the protection policy. For example, **File System Prod Daily Backups**.
 - c. In the **Type** field, select **File System**.
 - d. Click **Next**.
The **Purpose** page appears.
4. To create local backup protection, click **Self-Service Protection**.
5. Click **Next**.
The **Assets** page appears.
6. 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 the   icons 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.

7. On the **Objectives** page, 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 Administration and User Guide* provides instructions.
8. If you selected **Self-Service Protection** in the **Purpose** page:
 - a. Click **Add** under **Primary Retention**.
The **Add Primary Retention** dialog appears.
 - b. In the **Retentions (Self Service)** pane, specify the retention period for the self-service backups.
By default, all backup types have the same retention time. To change the retention times for specific backup types, for a File System agent installed in PowerProtect Data Manager 19.9 or later, clear **Set the same retention time for all backup types** and change the **Retain <backup_type> For** field values as required.
① **NOTE:** Copies created for self-service policy by a File System agent installed in PowerProtect Data Manager 19.8 or earlier, will continue to have the same retention for full and synthetic full backups.
9. On the **Target** pane of the **Add Primary Retention** dialog, specify the following fields:
 - a. **Storage Name**—Select a backup destination from the list of existing DD systems, or select **Add** to add a system and complete the details in the **Storage Target** window.
① **NOTE:** The **Space** field indicates the total amount of space, and the percentage of available space, on the storage system.
 - b. **Storage Unit**—Select whether this protection policy should use a **New** storage unit on the selected DD 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-ppdm-daily-123ab (300 GB/1 TB)**.
When you select **New**, a new storage unit in the format *policy name host name unique identifier* is created in the storage system upon policy completion. For example, **testvmpolicy-ppdm-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 on the selected system. PowerProtect Data Manager uses Governance mode for retention locking, which means that the lock can be reverted at any time if necessary. Toggling the **Retention Lock** slider on or off applies to the current backup copy only, and does not impact the retention lock setting for existing backup copies.

NOTE: If you enable **Retention Lock** for a replicated backup, set the **Retain For** field in the **Add Replication** backup schedule dialog to a period that is at least as long as the retention lock period for the primary backup, so that the replicated backup does not expire before the primary backup.

- e. **SLA**—Select an existing service level agreement that you want to apply to this schedule from the list, or select **Add** to create an SLA within the **Add Service Level Agreement** wizard.

The *PowerProtect Data Manager Administration and User Guide* provides instructions.

10. Click **Save** to save your changes and return to the **Objectives** page.

The **Objectives** page updates to display the name and location of the target storage system under **Primary Retention**.

NOTE: After completing primary retention settings, you can change any details by clicking **Edit** next to the primary retention stage.

11. Optionally, extend the retention period for the latest primary retention copy:

Extended retention provides more information about **Extend Retention** functionality.

- a. Click **Extend Retention** next to **Primary Retention**. An entry for **Extend Retention** is created below **Primary Retention**.

b. Under **Extend Retention**, click **Add**. The **Add Extended Retention** dialog appears.

c. **Retain the next scheduled full copy every...**—Specify a weekly, monthly, or yearly recurrence for the extended retention backup schedule.

d. **Repeat on**—Depending on the frequency of the full backup schedule, specify the day of the week, the date of the month, or the date of the year that the extended retention backup occurs.

e. **Retain For**—Specify the retention period for the backup. You can retain an extended retention backup for a maximum of 70 years.

f. Click **Save** to save your changes and return to the **Objectives** page.

12. Optionally, replicate these backups to a remote storage system:

a. Click **Replicate** next to **Primary Retention** or **Extended Retention**. An entry for **Replicate** is created to the right of the primary or extended retention.

NOTE:

PowerProtect Data Manager supports replicating an extended retention only if the primary retention already has one or more replication stages. Also, for replication of an extended retention, you can only select from the DDs to which the primary stage replicates.

For example, if there are 6 DDs available (DD001-DD006), and the primary backup is on DD0001:

- Replicate1, which is based on the primary retention, replicates to DD002.
- Replicate2, which is based on the primary retention, replicates to DD003.
- Extended retention is backed up to DD001.
- Replicate3, which is based on the extended retention, must replicate to DD002 or DD003.

b. Under **Replicate**, click **Add**. The **Add Replication** dialog appears.

c. Complete the schedule details in the **Add Replication** dialog.

NOTE: The schedule frequency can be every day, week, month, or x hours for replication of the primary backup, and every day, week, month, year, or x hours for replication of the extended retention backup. 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 1 to 12 hours.

All replication copies of the primary backup schedule use the same retention period, and by default, this retention period is inherited from the **Retain For** value of the synthetic full backup schedule. To specify a different retention period for all of the replication copies of this primary backup schedule, click **Edit**, change the value in the **Retain For** field, and then click **Save**. This retention period will be applied to all of the replicated copies (synthetic full and full) of this primary backup schedule.

13. Optionally, add a cloud stage for a primary, replication, or extended retention schedule to move backups from DD storage to the cloud tier:

NOTE: To move a backup or replica to the cloud tier, schedules must have a retention time of 14 days or more. PowerProtect Data Manager also requires the discovery of a DD system with a configured cloud unit.

- a. Click **Cloud Tier** next to **Primary Retention** or **Extend Retention** or, if adding a cloud stage for a replication schedule that you have added, click **Cloud Tier** under **Replicate**. An entry for **Cloud Tier** is created to the right of the primary or extended retention, or below the replication schedule.
- b. Under the entry for **Cloud Tier**, click **Add**.
The **Add Cloud Tier Backup** dialog appears, with summary schedule information for the parent node to indicate whether you are adding this cloud tier stage for the primary retention, the extended retention, or the replication schedule.
- c. Complete the schedule details in the **Add Cloud Tier Backup** dialog, and then click **Save** to save your changes and return to the **Objectives** page.
The *PowerProtect Data Manager Administration and User Guide* provides detailed instructions for adding a cloud stage for a primary, replication, or extended retention schedule.

14. Click **Next**.

The **Options** page appears.

15. On the **Options** page, select the additional option if required for the policy:

- **Troubleshooting**—Select this option to enable the debug logs for troubleshooting at higher debug levels. 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: Overriding the bug level in this manner may result in larger logs that may slow backup operations.

In Windows environments, the impacted logs include:

- o `FSAgentInstallPath\logs\vsacr.log`
- o `FSAgentInstallPath\logs\nsriscsi.log`
- o `FSAgentInstallPath\logs\nsriscsi_***.log`
- o `FSAgentInstallPath\logs\nsrwriter.log`
- o `FSAgentInstallPath\logs\ddfscn.***.log`
- o `FSAgentInstallPath\logs\ddfscn_***.log`
- o `FSAgentInstallPath\logs\ddfsv.log`
- o `FSAgentInstallPath\logs\ddfsv_***.log`
- o `FSAgentInstallPath\logs\ddfsrc_***.log`

In Linux environments, the impacted logs include:

- o `/opt/dpsapps/fsagent/logs/nriscsi.log`
- o `/opt/dpsapps/fsagent/logs/ddfscn.***.log`
- o `/opt/dpsapps/fsagent/logs/ddfsv.log`

NOTE: 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.

16. Click **Next**.

The **Summary** page appears.

17. 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.

When the new protection policy is created and assets are added to the protection policy, PowerProtect Data Manager performs backups according to the backup schedule.

18. 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.

You can monitor and view detailed information in the **Jobs** window for self-service backups and restores of database application agents.

NOTE: The **Cancel** and **Retry** options are not available for self-service jobs that are created by database application agents.

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 PowerProtect DD.


Add a policy to exclude assets from data protection operations



An exclusion policy prevents the File System agent from including certain assets in a protection policy, for example, assets that you intend to back up separately. The process of adding a policy is similar for all policy types. However, these instructions contain only elements and options that appear when you select the exclusion policy type.

Prerequisites

Review the prerequisites in Adding a protection policy for File System protection.

Steps

1. Select **Protection > Protection Policies**.
The **Protection Policy** window appears.
2. Click **Add**.
The **Add Policy** window appears.
3. In the **Type** page, specify the new protection policies group fields. For example, if you are creating a protection policy for daily backups in the Windows 2012 Server:
 - a. In the **Name** field, specify the name of the protection policy. For example, **File System Prod**.
 **NOTE:** The name that you specify here becomes part of the DD MTree entry.
 - b. In the **Description** field, specify a short description of the protection policy. For example, **File System Prod Daily Backups**.
 - c. In the **Type** field, select **File System**.
 - d. Click **Next**.
The **Purpose** page appears.
4. To exclude assets within the protection policy from data protection operations, click **Exclusion**.
5. Click **Next**.
The **Assets** page appears.
6. 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 the   icons 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.
7. Click **Next**.
The **Summary** page appears.
8. 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.

When the new protection policy is created and assets are added to the protection policy, PowerProtect Data Manager performs backups according to the backup schedule.
9. 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.

Cancel a File System agent backup or restore job

Starting with PowerProtect Data Manager 19.9, 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 Canceled 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.

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.
 - **Extended Retention**. If you choose this type, go to step 6.
 - **Replication**. If you choose this type, go to step 7.
 - **Cloud Tier**. If you choose this type, go to step 8.

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.
 - 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, and 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.
 - e. If applicable, select **Retention Time Objective**, and specify the number of Days, Months, Weeks, or Years.

NOTE: For compliance validation to pass, the value set for the **Retention Time Objective** must match the lowest retention value set for the backup levels of this policy's target objectives. For example, if you set the synthetic full backup **Retain For** to 30 days but set the full backup **Retain For** to 60 days, the Retention Time Objective must be set to the lower value, in this case, 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 **Extended Retention**, specify the following fields regarding the purpose of the new Extended Retention 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, and indicates the maximum targeted period in which data (transactions) might be lost from an IT service due to a major incident.
- NOTE:** By default, the RPO provides a grace period of 1 day for SLA compliance verification. For example, with a weekly extended retention schedule, PowerProtect Data Manager provides 8 days for the RPO to pass the SLA Compliance verification.
- 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.
- d. If applicable, select **Retention Time Objective**, and specify the number of Days, Months, Weeks, or Years.
- 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 **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.
- d. If applicable, select **Retention Time Objective**, and specify the number of Days, Months, Weeks, or Years.
NOTE: For compliance validation to pass, the value set for the **Retention Time Objective** must match the lowest retention value set for the backup levels of this policy's target objectives.
- 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.
8. 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: For compliance validation to pass, the value set for the **Retention Time Objective** must match the lowest retention value set for the backup levels of this policy's target objectives.
- d. If applicable, select the **Verify Retention Lock is enabled for all copies** option. This option is disabled by default.
- e. Click **Finish**.
9. 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**.
10. In the **Objectives** row of the **Summary** window, click **Edit**.
11. 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.
12. Click **Finish**.
 An informational message appears to confirm that PowerProtect Data Manager has saved the protection policy.
13. 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 interval of time between the most recent backup of the asset and the compliance check is greater than the RPO requirement, then an alert indicates 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 interval of time between the most recent backup of the asset and the compliance check is less than the RPO requirement, the RPO of the asset is in compliance.

If multiple assets in a policy are out of compliance at the same time when a compliance check occurs, a single alert is generated and includes information for all assets that are out of compliance 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.

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

Extended retention

You can extend the retention period for the primary backup copy for long term retention. For example, your regular schedule for daily backups can use a retention period of 30 days, but you can extend the retention period to keep the 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. When you extend the retention period of a backup in a protection policy, you 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 stage for a protection policy, you define a set of 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 an ad-hoc 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 schedule configuration.

When you configure extended retention to match backups by weekday, PowerProtect Data Manager may identify a backup that was taken on one weekday as being taken on a different 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 that 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 that 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 that 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, automatic extended retention creates a job at the beginning of the backup window for the primary stage. This job remains queued until the end of the backup window and then starts.

The following examples describe the behavior of backups with extended retention for centralized and self-service protection.

Centralized protection

For an hourly primary backup schedule that starts on Sunday at 8:00 p.m. and ends on Monday at 6:00 p.m., with a weekly extended retention schedule that is 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.

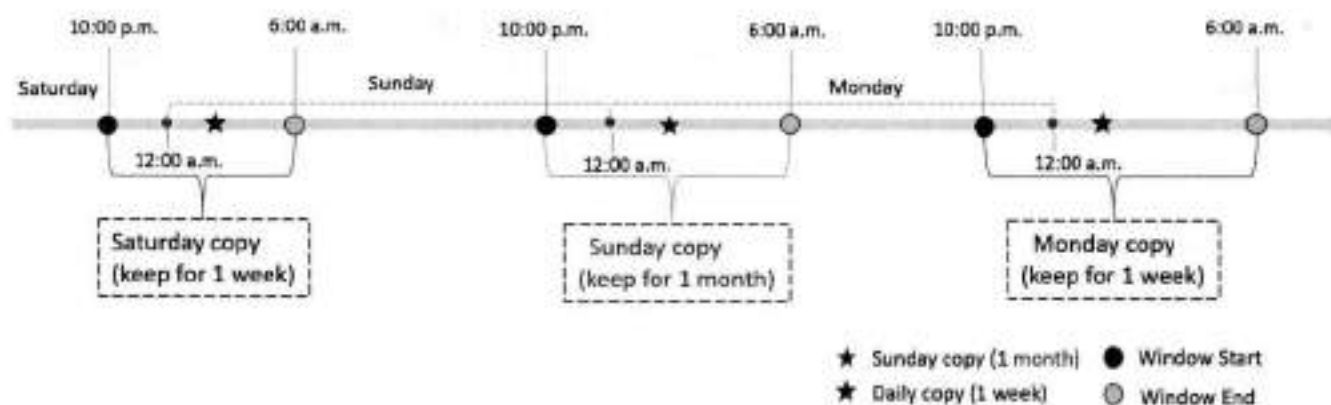


Figure 1. Extend retention backup behavior

Self-service protection

For self-service backups, PowerProtect Data Manager uses a default backup window of 24 hours. For a backup schedule that starts on Sunday at 12:00 p.m. and ends on Monday at 12:00 p.m. with a weekly extended retention schedule that is 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 stage by adding a replication stage to the extended retention backup. The rules in the extended retention stage 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 stage for the primary backup.
- Configure the replication stage of the extended retention and match this stage with one of the existing replication stages based on the primary backup. Any changes to a new or existing storage unit in the extended retention replication stage or the replication stage of the primary backup is applied to both replication stages.
- The replication stage of extended retention backups only updates the retention time of replicated backup copies and does not create any new backup copies 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**.
2. 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.
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 the storage icon to the right of the icon for the asset, for example, **DD**. The table in the right pane lists the backup copies.
5. Select one or more backup copies from the table, and click **Edit Retention**.
6. 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.

7. 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 upon expiration of the retention period, 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

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 the storage icon to the right of the icon for the asset, for example, **DD**. The table in the right pane lists the backup copies.
5. 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.
6. 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**.
7. 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 is 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 is shown in UTC. Go to **Alerts > Audit Logs** to view the audit log.

8. 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

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 the storage icon to the right of the icon for the asset, for example, **DD**. 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 **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 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 successfully deleted from protection storage. If the deletion is successful, the deleted copies no longer appear in the table.

Export data for deleted 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.
-  **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 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 the storage icon to the right of the icon for the asset, for example, **DD**. 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 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.

Exclusion filters

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 and to backups taken through self-service CLI.

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 or smaller than a specified size. Specify a value in either the Greater than or Less than 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.

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.

- When you are finished building the filter, click **Add Filters**. The new filter appears in the table.
- You can add up to four filters using the previous steps. When you are finished, click **Next**.
- In the **Summary confirmation** page, verify the filter information and click **Finish**.

Guidelines for exclusion filters

Review the following guidelines for exclusion filters.

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.

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. For example, C:\\abc*.txt matches all text files in the folder C:\\abc. All matching files under subfolders of that specific 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 **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 7. 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.

Table 7. Absolute and relative path examples (continued)

Type of path	Folder	File
Relative	*\\folder1\\folder2*	*\\folder1\\folder2\\sample.log
	In this example, the filter excludes all files and folders under any volume with the hierarchy folder1\folder2.	In this example, the filter excludes all sample.log files under any volume with the hierarchy folder1\folder2.
Absolute	D:*\\folder1\\folder2*	D:*\\folder1\\folder2\\sample.log
	In this example, the filter excludes all files and folders under any folder in D: with the hierarchy folder1\folder2.	In this example, the filter excludes all sample.log files under any folder in D: with the hierarchy folder1\folder2.

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 the filtering conditions.

Steps

1. Select **Protection > Filters**.

The **Exclusion Filters** window appears, which displays the following information:

- Name
- Description
- Conditions
- Logical Operator

2. To edit a filter, complete the following tasks:

- a. Select a filter, and click **Edit**.
The **Edit Exclusion Filter** wizard appears.
- b. Modify the desired fields, and then click **Next**.
The **Summary** 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 already exist.

About this task

To create a protection policy for file system protection and apply an exclusion filter to it, follow the steps in *Add a policy to exclude assets from data protection operations*.

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 **Summary** 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

After the backup starts, you can view details about the files that are excluded from the protection policy. To view the excluded files:

- Open the **Jobs** window and select the job.
- Click the **Details** icon to the left of the job name.
- In the **Task Summary** section, click the link that indicates the total number of tasks.
- Click the **Details** icon to the left of the task, and then review the protection policy details and excluded files.

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 **Summary** page appears.
3. Select **File Exclusions > Edit**.
4. Clear the checkbox 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/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 the  icons 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 checkbox next to the desired file system and click **View Copies**.

You can also use the **Search** field, 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 **DD**, and then select one of the available copies that appear in the table.
4. In the right pane, select the checkbox 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.

i **NOTE:** If the destination file system asset already contains some data, this data will be overwritten.

6. On the **Summary** page:
 - a. Review the information to ensure that the restore details are correct.
 - b. Click **Restore**.
7. 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_ybind 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.

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 the   icons 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 checkbox next to the file system and 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.

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 **DD**, and then select from one of the available copies that appear in the table.
4. In the right pane, select the checkbox 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 same or original machine.
 - Restore to alternate host.
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.
 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.

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. On the **Summary** page:
 - a. Review the information to ensure that the restore details are correct.
 - b. Click **Finish**.
 10. Go to the **Jobs** window to monitor the restore.
A restore job appears with a progress bar and start time.

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/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.

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.

Enable the File System agent after hostname change

After the hostname of the File System agent host is changed, you must update the lockbox setting for the protection policy.

About this task

Perform the following steps to enable the File System agent operations after the hostname is changed.

Steps

1. Re-register the File System agent with PowerProtect Data Manager:
 - For Linux, run the `install.sh` script from the `agentavc` directory. For more information, see [Install the File System agent on Linux](#).
 - For Windows, run the `fsagent-19.10.0.0.exe` program, and select **Change** when prompted for the action to perform. For more information, see [Install the File System agent on Windows in interactive mode](#).
2. 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.
 - NOTE:** If the File System agent is installed to a non-default path, delete all of the files in the `lockbox` subdirectory of the installation directory.
3. In the PowerProtect Data Manager UI, configure the lockbox:
 - a. In the left navigation pane, select **Protection > Protection Policies**.
 - b. On the **Protection Policies** page, select the applicable protection policy in the list and click **Set Lockbox**.
4. From the protection policy, remove any assets that were protected under the old hostname.
 - NOTE:** Until steps 4 through 6 are followed, any asset that was protected under the old hostname is no longer protected under the new hostname.
5. Run a manual discovery.
6. Add any asset now in an available state back to the protection policy.

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 using the centralized PowerProtect Data Manager restore functionality, or directly 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 ddfsadmin 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 will display a list of local backups in C:\ taken within the last five days.

ddfsadmin sync

This command ensures 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 ddfsadmin command. Note that 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 `ddfsadmin 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 `ddfsadmin` utility for file systems provides more information about the `ddfsadmin 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
```

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 taken in the last 55 days.

```
[root@XXXX ~]# ddfsadmin 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.

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 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
- Performing a bare-metal recovery of Windows clusters
- Performing application restores after bare-metal recovery

Disaster recovery limitations

The following limitations apply to performing disaster recovery in the File System agent.

- Distributed File System (DFS) is not supported by disaster recoveries.
- Disaster recovery data of clusters does not include critical disks in shared storage. Critical disks must be backed up separately.

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 non-critical volume that has a critical volume mounted on it, or any non-critical 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 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 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 checkbox 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 non-critical 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 non-critical 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 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 file. 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 in a stable state.

- When performing an SSR, ensure that there is at least 50% of free space on the system disk.

About this task

Steps

1. In PowerProtect Data Manager, from the left pane select **Restore > Assets**.
2. Select the checkbox 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 a storage system to display the copies on that storage system.
5. Select the desired backup copy and 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 (DSRM) provides more information.
2. Recover the Active Directory. Restore Active Directory from the disaster recovery backup provides more information.
3. Perform an authoritative or nonauthoritative Microsoft restore based on user configuration. Authoritative and nonauthoritative restore 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 `msconfig` 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 the storage targets 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.

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. In addition, 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 deselect **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
 - i** **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)

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 a 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 the 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, some unallocated space might remain. You can extend the partition size after the BMR 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.
 - 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, 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 a BMR

About this task

Download, modify, and deploy the custom WinPE image for a 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, and then add the drivers to the WinPE image. For more information, see Add NIC or disk drivers to the custom WinPE image.

- 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

- 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.
- Select **Settings > Downloads**.
- Select **WinPE**.
- Click **Download** to download the .iso file of the custom WinPE image.
- 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.

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

- Open the .iso file of the WinPE image with a utility like UltraISO or MagicISO.
- 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.

4. 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:** WDS is only one method of booting from a WinPE image over the network, but other boot methods are unsupported.


Steps

1. Configure the WDS server.
2. Add the WinPE image to the boot menu.
3. Ensure that PXE booting is enabled on the WDS target host.
4. 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

1. 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.
2. 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. 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 either the **IPv4** tab to configure the network to communicate with PowerProtect Data Manager during the BMR.
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 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 click **Next**.
- NOTE:** In **Server Name or IP**, enter only the FQDN or IP of the server.

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 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 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: 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, click **OK**.

- e. Perform one of the following actions:

- To cancel the BMR, click **Cancel**.
- To proceed with the BMR, click **Restore**.

16. If you clicked **Restore**, wait until the BMR 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 is complete:


- The final status of the BMR 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 `ddfsvc` 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 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\Essagent\logs`.
4. Copy the files or folder to the mounted shared drive by running the following command:

```
copy ddsrc.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 a 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, the 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 `ddsrc.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.

Performing a bare-metal recovery of Windows clusters

Bare-metal recovery (BMR) can restore a Windows cluster configuration, but you must ensure that the shared disk data and user data are protected using the relevant application agents. BMR restores only the system state and critical disks on the cluster nodes.

After the BMR, use normal File System agent backups to restore file system data, non-critical disks, and critical disks on shared storage. Also, use application-agent backups to restore application data.

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:

- Installation and operation
- Backups
- Disaster recovery
- Restores
- Storage units

Installation and operation

You might encounter the following issues while installing or operating the File System agent.

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.

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.

4. 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.

5. Start the agent service.

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.emc.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.

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\DPSPAGENT\settings\ddfssv.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.

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 `ddfsrv` error

Backing up BMR data can fail with the error `ddfsrv 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.

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:

```
ddfsrv Error <0000>: Unable to select W component from V for restore: The specified object was not found. (VSS error 0x80042308)
```

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 `pendingfilereoperations` 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`, `ddfarc.exe`, `ddfacon.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 are not deleted after a failed or canceled recovery operation

After a failed or canceled recovery operation, you might see files continuing to exist in `C:\dbapps_temp_dir`. Replace `C:` with the drive to which Windows was installed.

These files cannot be manually deleted. To remove the files, contact Customer Support.

Restores

You might encounter the following issues while performing restores with the File System agent.

Centralized file-level restore operations

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 to `G:\CFLR\D\Folder\file1.txt` instead of `G:\CFLR\file1.txt`.

There is no workaround.

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 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 a network connectivity error or an agent host timed out error

You might see an error similar to one of the following when attempting to restore a clustered drive:

```
ARA0005: Unable to restore FILE_SYSTEM asset C:\ClusterStorage\Volume3 because of a network connectivity issue on agent host v2019c2.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.
```

```
The File System restore was unsuccessful because the request to the Agent host timed out.
```

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 Administration and User 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.

Storage units

You might encounter the following issues with storage units:

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.

Dell RecoverPoint for Virtual Machines

Enterprise grade, hypervisor level data protection for VMware virtual machines

Summary

- Protects VMware Virtual Machines with VM level granularity
- vAdmins work from VMware vCenter via a plug-in
- Supports all storage and application types

Value

- Streamline data protection workflows with reliable and repeatable processes
- Respond faster to changing business and data protection needs
- Shorten application development cycles by providing a replica for isolated test and development use
- Enable datacenter migration with minimal interruption
- Leverage offsite replication for backup operations with no impact to the production site
- Empower vAdministrators to meet required data protection service level agreements (SLAs)
- Integrate with VMware vRealize Operations Manager via Dell Storage Analytics (ESA) with deep visibility into your virtualized infrastructure

Dell RecoverPoint for Virtual Machines redefines data protection for VMware Virtual Machines (VMs) enabling local, remote, and concurrent local and remote replication with continuous data protection for on-premises recovery to any point-in time (PIT).

VMware hypervisor-based, the solution is storage and application agnostic, with built-in orchestration and automation accessible via VMware vCenter web client plug-in.

Gain control over your VMware environment

- Enable continuous data protection for on-premises any PIT recovery for near zero RPO and RTO.
- Ensure recovery consistency for interdependent applications.
- Provide synchronous (sync) or asynchronous (async) replication policies.
- Protect data using proprietary Consistency Groups (CG) and Group Sets, ensuring recovery consistency for one application or interdependent applications.
- Manage blueprints direct from VMware vRealize Automation (vRA).
- Support multiple sites with up to 4:1 fan-in for centralized DR site protecting multiple branch offices and 1:4 fan-out replication for recovery, test and development operations.



vCenter integration

With its tight VMware integration, RecoverPoint for Virtual Machines protect VMs with VM-level granularity. Its vCenter plug-in empowers vAdministrators to protect single or multiple VMs locally or remotely to the target site as well as to perform automated discovery, provisioning and orchestration for disaster recovery testing.

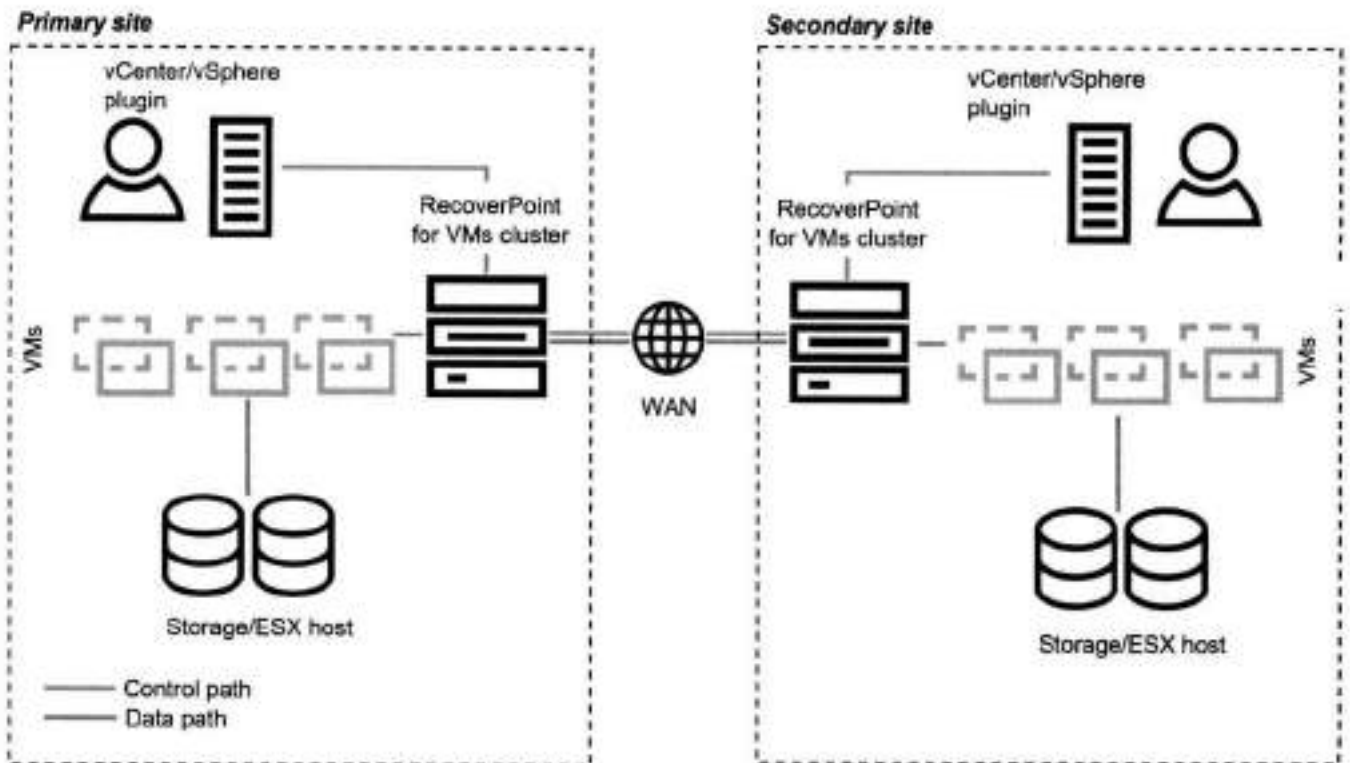



Figure 1: Machine-level granularity for any-point-in-time recovery of VMware environments.

Take the next step

Contact your Dell sales representative or authorized reseller to learn more about how RecoverPoint for Virtual Machines and Data Protection Suite can benefit your organization.



[Learn More](#) about
RecoverPoint for
Virtual Machines



[Contact](#) a Dell Technologies Expert

Dell PowerProtect Data Manager

Next generation software platform for proven and modern cloud data protection

Essentials

- **Software-defined data protection**
- **Autonomous operations:** Automated discovery and protection of databases, virtual machines, file systems and Kubernetes containers
- **Multicloud optimized:** Extend protection with backup to cloud, backup in-cloud, long term retention and cloud disaster recovery
- **Unique VMware protection:** Ensure availability of all your VMs at scale without business disruption
- **Cyber recovery:** Increase business resiliency to rapidly recovery from cyber incidents
- **Self-service backup and restore:** Enable data owners from their native interfaces
- **Centralized oversight and governance:** Mitigate risk and assures compliance of SLAs and SLOs
- **Simple protection workflows:** Minimize daily operations
- **Cloud-based monitoring and analytics**
- **Efficient protection:** Protect data directly to PowerProtect appliances
- **PowerStore backup and recovery support**
- **PowerProtect Data Manager Appliance:** Integrated data protection platform in a single appliance

Discover, manage, protect and restore data

- Kubernetes
- VMware, Hyper-V and open hypervisors
- Oracle, Microsoft SQL and Exchange, SAP HANA
- Windows and Linux Filesystems

Gain the confidence that your data is protected and available to drive value as a business asset

The IT landscape has changed. Reasons to protect workloads extend beyond IT-driven application restores and disaster recovery scenarios. Backup requirements have transcended IT teams and have crossed over to application and data owners who aspire to do more than simply restore their data.

To support these expanding use cases and requirements, backup applications are transforming to provide more than just access to backups and restore capabilities including:

- Analysis and reuse for dev/test
- Leverage the cloud to extend data center capabilities
- Protect cloud native applications
- Enable self-service backup and restore from native applications
- Maintain centralized governance and control
- Increase business resiliency to rapidly recovery from cyber incidents

To address these requirements, Dell PowerProtect Data Manager is at the forefront of this transformation to modern data protection.



Next generation cloud data protection

You can address these challenges by transitioning to a broader cloud data protection strategy. One that places you on the path to modernizing your data center and unlocks the value of your data and applications for everyone. This is an evolution – and it begins with PowerProtect Data Manager.

Data Manager gives you valuable insight into protected on-premises and in-cloud workloads, applications, file systems, and virtual machines (VMs). Plus, complete oversight and governance to ensure compliance.

Designed with operational simplicity and agility in mind, Data Manager enables the protection of traditional workloads including Oracle, Exchange, SQL, SAP HANA and file systems as well as Kubernetes containers and virtual environments. Restore data on-premises or in the cloud. Governance control ensures IT compliance, making even the strictest service level objectives obtainable.





Autonomous operations

Automatically discover and protect databases, VMs, file systems and Kubernetes containers, while a common policy engine automates compliance and governance across workloads. You can instantly access protected VM images to support new use cases such as quickly deploying development and test environments. Data Manager integrations provide native vSphere Storage Policy Based Management integration for VM protection, offering storage and backup admins, as well as VM owners, the ability to choose a storage policy to apply to every VM automatically when it is instantiated.

Multicloud optimized

Leverage the cloud for backup, long-term retention and disaster recovery. Whether you're focusing on private, public or hybrid cloud, you can be confident your data is protected at the level you need.

Data Manager extends protection to the cloud by tiering backups to cloud storage for long-term retention to minimize costs and maximize access to backups without impacting on-premises protection storage resources.


Data Manager protects in-cloud workloads running on AWS, Azure and Google Cloud, as well as enables disaster recovery to the public cloud. Automated orchestration of fail-over, failback and testing simplify production disaster recovery scenarios. Easy to install and deploy from the AWS, Azure and Google Cloud Marketplaces, Data Manager along with PowerProtect DD Virtual Edition deliver a high level of performance and efficiency through deduplication.

Data Manager also brings enterprise data protection for the VMware Tanzu portfolio, both on-premises and in the cloud. With VMware running Kubernetes everywhere, enabling the protection of Tanzu is essential for business operations.

Protect cloud-native workloads across multiple public clouds

Mission and business-critical applications deployed in public clouds require cloud-native methods to protect their data. Unfortunately, the level of native data protection available in public clouds isn't sufficient, consistent nor designed to reign in sprawl.

Use a single tool to discover, orchestrate and automate the protection of AWS and Azure workloads via powerful tag-based policies and REST APIs. As a SaaS component of Data Manager, PowerProtect Cloud Snapshot Manager protects cloud-native workloads across multiple public clouds. This provides you with global visibility and control and enables you to gain insight into data protection activities across your public cloud infrastructure.



Change the way you protect VMware VMs with Transparent Snapshots

The volume of VMware data continues to grow and protecting that data at scale will only become more challenging. While alternate approaches have attempted to overcome the issues of VM latency and business disruption, all are fraught with undesirable compromises around latency, cost, scalability, performance and complexity. Transparent Snapshots enables you to protect your VMware environments more effectively while overcoming these challenges.

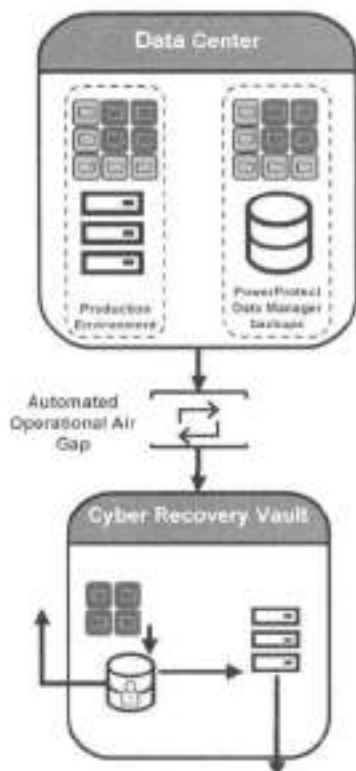
Transparent Snapshots simplifies and automates VM image-level protection and backs up virtual machines without the need to pause them during the backup process. The result is significantly reduced impact to business operations, especially on large, high-change-rate VMs. The simplified backup process also reduces infrastructure costs by removing the reliance on proxies for data movement.

Transparent Snapshots delivers up to 5x faster backups¹, up to 6x faster restores² and up to 5x reduction in VM latency³, effectively and efficiently backing up your VMs via a process that requires fewer steps. The result is less impact to your entire VMware environment, ensuring availability of all your VMs without business disruption.

Increase business resiliency with cyber recovery capabilities

Protecting your business starts with protecting your data. To reduce business risk caused by cyber attacks and to create a more cyber resilient approach to data protection, you can modernize and automate your recovery and business continuity strategies and leverage the latest intelligent tools to detect and defend against cyber threats.


As part of Data Manager, PowerProtect Cyber Recovery provides proven, modern and intelligent protection to isolate critical data, identify suspicious activity and accelerate data recovery allowing you to quickly resume normal business operations.



Self-service for data owners combined with central IT governance

Extend data protection for expanding use cases while maintaining control by empowering data and application owners to perform self-service backup and restore operations from native applications directly to Dell PowerProtect appliances. At the same time, Data Manager provides IT with the necessary oversight and governance to ensure compliance.

Data owners and administrators are also empowered with cloud-based monitoring and analytics through Dell CloudIQ. Cloud IQ provides telemetry, machine learning and predictive analytics to proactively take action and speed time to resolution.



Fast, simple backup and recovery using Storage Direct with PowerStore

PowerProtect Data Manager can also be leveraged in the same environment as your PowerStore and PowerProtect appliances to provide centralized management and orchestration for backups and recoveries. With all the benefits of Data Manager and PowerProtect DD, including Instant Access, image-level restores, retention lock support and Transparent Snapshots, you can easily protect multiple PowerStore clusters with crash-consistent snapshots.


Start your journey towards modern data protection with the PowerProtect Data Manager Appliance

The PowerProtect Data Manager Appliance includes PowerProtect Data Manager to deliver next generation multicloud data protection. It offers complete backup, deduplication, replication, recovery, instant access and restore, search and analytics, and seamless VMware integration – plus, cloud readiness with disaster recovery, long-term retention to the cloud as well as support for multicloud workloads. It is easy to deploy, manage and upgrade, and grows in place from 12TB to 96TB in 12TB increments – all with license keys, requiring no additional hardware, downtime, or complexity.⁴

To learn more about the PowerProtect Data Manager Appliance, visit the [webpage](#) today.

Take the next step

Contact your Dell sales representative or authorized reseller to learn more about how PowerProtect Data Manager can benefit your organization.



[Learn More](#) about
PowerProtect Data
Manager



[Contact](#) a Dell Technologies Expert

¹ When comparing PowerProtect Data Manager 15.13 with Transparent Snapshots backup performance to PowerProtect Data Manager with VADP backup performance. Based on Dell internal testing, June 2023.

² When comparing PowerProtect Data Manager 15.13 with Transparent Snapshots restore performance to PowerProtect Data Manager with VADP restore performance. Based on Dell internal testing, June 2023.

³ When comparing PowerProtect Data Manager 15.13 with Transparent Snapshots VM latency performance to PowerProtect Data Manager with VADP VM latency performance. Based on Dell internal testing, June 2023.

⁴ If starting at less than 24TB, a single field upgrade kit is required to bridge from 24TB to larger capacities.



EQUIPAMENTOS DELL POWERPROTECT SÉRIE DD

Principais benefícios

Rápido, seguro e eficiente

- 1,5 PB de capacidade útil em um rack
- Até 3 PB de capacidade para retenção em longo prazo
- Capacidade lógica de até 30% por TB*
- Recursos instantâneos de análise e restauração de até 54 VMs e 100.000 iOPS****
- Conectividade de rede de alta velocidade – 10 GbE, 25 GbE e 100 GbE
- Integração perfeita e desempenho superior com o PowerProtect Data Manager e o Data Protection Suite
- Suporte aos principais aplicativos empresariais de backup e arquivamento

Proteção multicloud líder no setor

- Armazenamento de proteção definido por software no local e na nuvem com o PowerProtect DD Virtual Edition (DDVE)
- O DDVE tem um dimensionamento de até 256 TB na nuvem
- Melhora o desempenho de restauração na nuvem em até dez vezes***
- O Cloud Tier fornece retenção em longo prazo simples e eficiente para nuvem pública, privada ou híbrida
- Recuperação de desastres de baixo custo para a nuvem

Simplicidade operacional

- O System Manager aprimorado oferece visualização completa do chassis
- Ponto único de gerenciamento para toda a série DD por meio do PowerProtect DD Management Center
- Compatibilidade com o Smart Scale, o que reduz a complexidade no gerenciamento de dados conforme a necessidade

Eficiência no uso de energia com a linha Dell Storage

- A Dell tem o compromisso de aumentar a eficiência no uso de energia a cada geração de novos portais de armazenamento.

O melhor equipamento para um armazenamento protegido

A série DD facilita que as organizações protejam, gerenciem e recuperem dados conforme a necessidade em seus diversos ambientes. A série DD é a última geração de equipamentos Data Domain, agora a referência em proteção de dados da borda ao núcleo e à nuvem. A série DD oferece a compatibilidade com ecossistemas, a eficiência, a proteção de dados avançada e os recursos prontos para a nuvem que os clientes esperam e gostam do Data Domain e os eleva a um novo patamar.

O DD Operating System (DDOS) é a inteligência que potencializa a série DD. Gera a agilidade, a segurança e a confiabilidade que permitem à série DD oferecer o armazenamento com proteção multicloud escalável, em alta velocidade e líder do setor para fazer backup, arquivamento e recuperação de desastres. O DDOS se integra com perfeição às infraestruturas atuais, o que facilita a usabilidade com os principais aplicativos de backup e arquivamento e oferece um desempenho superior quando é usado com o Dell PowerProtect Data Manager e o Data Protection Suite. Ao comprar um equipamento da série DD novo, agora você pode consumir o DDOS na forma de assinatura, o que dá flexibilidade na implementação e diminui os custos iniciais.

Proteção de dados rápida, segura e eficiente

A série DD reduz o risco de perda de dados e aproveita o valor dos dados protegidos, cumprindo com Acordos de Nível de Serviço cada vez mais exigentes e aumentando o ROI. O DDOS faz a série DD oferecer backups até 38% mais rápidos e restaurações até 45% mais rápidas em níveis mais altos de compactação.** Esse nível maior de eficiência de compactação, com frequência, aumenta a capacidade lógica em até 30% por TB*.

Agora, a série DD pode ser dimensionada para mais em uma capacidade física de 1,5 PB em apenas um rack, assim utilizando o mínimo de espaço físico e diminuindo a energia e a refrigeração em até 41%***. Ao utilizar unidades de disco mais densas, a série DD reduziu em até 39% o espaço necessário em rack.

A série DD oferece até 3 PB de capacidade na nuvem para ter retenção em longo prazo, com o Cloud Tier.

A série DD comporta alta disponibilidade em apenas um rack. Com isso, ela pode diminuir ainda mais o custo total de propriedade reduzindo o tempo de inatividade no caso inesperado de uma falha de hardware. Também oferece conectividade de rede em alta velocidade compatível com adaptadores de rede de 25 GbE e 100 GbE.

* Com base em testes internos da Dell e dados de telemetria em campo, Março de 2022. Os resultados reais podem variar.

** Com base em testes internos da Dell em comparação com a geração anterior, março de 2022. Os resultados reais podem variar.

*** Ao comparar 1 petabyte de dados em um DD9880 com o Cloud Tier e no PowerProtect DD9900 com o Cloud Tier. Os resultados reais podem variar. Março de 2022.

**** Com base em testes internos da Dell ao comparar o DDVE 7.7 com o DDVE 7.1. Os resultados reais podem variar. Março de 2022.

***** Ao usar o DDOS 7.7 e posterior ao DD9800. Com base em testes internos da Dell. Os resultados reais podem variar. Março de 2022.

Smart Scale para equipamentos PowerProtect

Muitas vezes, as organizações devem gerenciar vários data centers e ambientes na nuvem, adicionar, fazer upgrade e desativar a infraestrutura de armazenamento de proteção, instalar novos aplicativos em constante evolução, além de otimizar a capacidade e o desempenho. Não é uma tarefa fácil, mas uma que a Dell está ajudando as empresas a superar com o Smart Scale. O Smart Scale permite gerenciar até 32 equipamentos da série DD dentro de um só pool de sistemas em um namespace unificado que reduz a complexidade do gerenciamento e aumenta a eficiência de armazenamento. O Smart Scale é implementado gratuitamente por meio de nosso console de gerenciamento centralizado, o PowerProtect DD Management Center. O Smart Scale é compatível com DD9900, DD9400, DD6900 e DD6400. Para integrar softwares, oferecemos a compatibilidade com Dell PowerProtect Data Manager, Dell NetWorker e aplicativos de backup de terceiros. O Smart Scale conta com unidades de armazenamento móveis que oferecem flexibilidade e mobilidade transparente dos dados de backup em cada pool.

Acesso instantâneo e restauração instantânea

Os recursos instantâneos de acesso e restauração oferecem alto desempenho de VMs com até 100.000 IOPS, com a capacidade de acessar instantaneamente até 64 VMs ao mesmo tempo.*****

Os recursos instantâneos de acesso e recuperação poupam tempo, o que reduz a média de tempo para reparo ao permitir o acesso instantâneo aos dados da imagem de backup nas unidades SSD da série DD incluídas. Eles também economizam espaço de armazenamento primário com a capacidade de gerenciar dados no próprio equipamento e reduzem os custos, utilizando melhor os recursos físicos tanto na proteção de dados quanto nos ambientes de produção.

Em caso de erro ou recuperação de desastres em um ambiente virtualizado, a série DD pode acionar VMs voltadas à produção imediatamente dentro do próprio equipamento. Ao fazer isso, o cliente pode continuar a rotina diária sem enfrentar tempo de inatividade, enquanto as VMs com falha são restauradas para o ambiente de produção.

Arquitetura de Invulnerabilidade de Dados

A série DD foi desenvolvida como o armazenamento de última instância, garantindo que você possa sempre recuperar seus dados com confiança. A Data Invulnerability Architecture é integrada ao DDOS e à série DD para oferecer a melhor defesa do setor contra a perda de dados. A verificação de leitura e gravação em linha fornece proteção e recuperação automática no caso de problemas de integridade dos dados durante a ingestão e a recuperação de dados, enquanto o RAID-6 e os hot spares protegem contra falhas de disco.

A captura e a correção de erros de E/S em linha durante o processo de backup eliminam a necessidade de repetir os backups, garantindo que eles sejam concluídos no prazo e atendam aos Acordos de Nível de Serviço. Além disso, diferentemente de outros arrays ou file systems empresariais, na série DD, a detecção contínua de falhas e a autocorreção garantem que os dados possam ser recuperados durante todo o ciclo de vida.



Verificação completa dos dados

Verificações completas de dados significam que será feita uma leitura dos dados após a gravação, e eles serão comparados ao que foi enviado ao disco para comprovar que estão acessíveis por meio do sistema de arquivos em disco e que não estão corrompidos. Especificamente, quando o DDOS recebe uma solicitação de gravação do software para backup, ele calcula um checksum dos dados. Depois de analisar os dados quanto à redundância, ele armazena os novos segmentos de dados e todos os checksums. Depois que todos os dados são gravados no disco, o DDOS verifica se pode ler todo o arquivo da bandeja do disco e pelo PowerProtect DD, além de verificar se os checksums dos dados lidos correspondem aos checksums dos dados gravados. Isso confirma que os dados estão corretos e podem ser recuperados de cada nível do sistema.

Portfólio abrangente da série DD

	DDVE 96 TB	DD3300	DD6400	DD6900	DD9400	DD9900
Inclusão de backup (com DD Boost)	Até 11,2 TB/h	Até 7,0 TB/h	Até 27,7 TB/h	Até 33 TB/h	Até 57 TB/h	Até 94 TB/h
Capacidade lógica (com nível ativo)	Até 4,8PB	Até 1,6PB	Até 11,2 PB	Até 18,7PB	Até 49,9PB	Até 97,5 PB
Capacidade útil (com nível ativo)	1 TB a 96 TB	4 TB a 32 TB	8 TB a 172 TB	24 TB a 266 TB	192 TB a 766 TB	576 TB a 1,5 PB

Capacidade lógica com base em deduplicação de até 50 vezes (DD3300) e deduplicação de até 85 vezes (DD6400, DD6900, DD9400 e DD9900) de acordo com a compactação adicional de dados assistida por hardware até 30% melhor do que a geração anterior. A capacidade e o throughput reais dependem da carga de trabalho de aplicativos, da deduplicação e de outras configurações.

Integração perfeita

A série DD se integra facilmente a infraestruturas existentes, viabilizando a usabilidade com os principais aplicativos de backup e arquivamento, e oferece desempenho superior quando combinada com o PowerProtect Data Manager e o Data Protection Suite.

A série DD aceita diversos métodos de acesso ao mesmo tempo, como NFS, CIFS, VTL, NDMP e DD Boost™, e todos os aplicativos e utilitários podem ser compatíveis com o mesmo equipamento da série DD para permitir maior consolidação do armazenamento de proteção. Um sistema pode se apresentar como servidor de arquivos, oferecendo acesso por NFS ou CIFS por Ethernet, como biblioteca de fitas virtuais (VTL) por Fibre Channel, como servidor de fitas NDMP por Ethernet ou como disco de destino usando interfaces específicas de aplicativo, como o DD Boost. O DD VTL é qualificado com os principais sistemas abertos e aplicativos de backup corporativo IBMi.

Proteção multicloud líder do setor

A série DD simplifica e tem eficiência operacional, incluindo resiliência e dimensionamento, à medida que você crescer em qualquer ambiente de nuvem, seja ela privada, pública ou híbrida. A série DD é compatível com o ecossistema de nuvem mais abrangente — AWS, Azure, VMware Cloud, Google Cloud, Alibaba Cloud e Dell ECS — para oferecer uma proteção de dados excelente na nuvem com custos reduzidos. A série DD pode armazenar nativamente em camadas os dados deduplicados em qualquer ambiente de nuvem compatível para oferecer retenção em longo prazo com o Cloud Tier. A série DD oferece recuperação de desastres rápida e orquestrada, além de uma arquitetura eficiente para estender a proteção de dados no local com custos reduzidos.

PowerProtect DD Virtual Edition

O PowerProtect DD Virtual Edition (DDVE) usa a capacidade do DDOS para oferecer armazenamento de proteção definido por software no local e na nuvem. É fácil e rápido fazer o download, configurar e implementar o DDVE, sendo possível começar a ser usado em poucos minutos. O DDVE pode ser implementado em qualquer hardware padrão, convergente ou hiperconvergente, e funciona no VMware vSphere, Microsoft Hyper-V e KVM, bem como na nuvem com AWS, AWS GovCloud, VMware Cloud, Azure, Azure Government Cloud, Alibaba Cloud e Google Cloud. Ele também é certificado com os servidores Dell PowerEdge e com o VxRail. Uma ferramenta de avaliação pode ser executada durante a implementação para verificar a infraestrutura subjacente e garantir que ela atenda aos requisitos recomendados. Uma instância só do DDVE pode ser dimensionada em até 256 TB na nuvem e até 96 TB no local. A capacidade pode ser facilmente transferida entre sistemas virtuais e/ou locais e pode ser dimensionada em incrementos de 1 TB, permitindo aumentar a capacidade conforme a demanda da empresa. O DDVE mantém os principais recursos do DDOS e inclui DD Boost, DD Encryption e DD Replicator. Ele pode ser configurado e gerenciado usando o DD System Manager e realiza o gerenciamento centralizado de várias instâncias do DDVE, no local e na nuvem, pelo PowerProtect DD Management Center.

Retenção em longo prazo e recuperação de desastres na nuvem

Com o Cloud Tier, o DDOS pode armazenar nativamente em camadas os dados em uma nuvem pública, privada ou híbrida para oferecer retenção em longo prazo. Somente dados exclusivos são enviados diretamente da série DD para a nuvem, e os dados chegam ao armazenamento em objeto na nuvem já deduplicados. É compatível com AWS, AWS Gov Cloud, Azure, Google Cloud, IBM Cloud, Alibaba Cloud, Seagate Lyve Cloud e Dell Elastic Cloud Storage (ECS). Com taxas de deduplicação de até 65 vezes, o espaço ocupado pelo armazenamento diminui significativamente, reduzindo o custo total de propriedade. O Cloud Tier pode ser dimensionado para até 3 PB de capacidade útil. Com DD Encryption, os dados na nuvem permanecem seguros. O Cloud Tier funciona com o DDVE para implementações no local.

O Cloud Disaster Recovery (Cloud DR) permite que as empresas copiem as VMs de backup dos ambientes da série DD no local à nuvem pública (AWS, VMware Cloud on AWS e Azure) e orquestram os testes de recuperação de desastres e failover de cargas de trabalho na nuvem em uma situação de desastre com orquestração completa.

Simplicidade operacional

A série DD conta com instalação e gerenciamento extremamente simples, resultando em custos administrativos e operacionais menores. Os administradores podem acessar o DDOS por meio da linha de comando via SSH ou por meio do DD System Manager, uma interface gráfica do usuário baseada em navegador.

Diversos equipamentos da série DD podem ser gerenciados e monitorados por meio de apenas uma interface, o PowerProtect DD Management Center, ou DDMC. Os painéis de indicadores com personalização dão visibilidade do status agregado e por localização, bem como a capacidade de se aprofundar nos dados do nível do sistema. Agora, o DDMC pode mostrar insights sobre as capacidades atuais e estimadas no nível do sistema para a série DD e os sistemas Data Domain preexistentes, o que proporciona mais previsão e gerenciamento da capacidade. O acesso baseado em funções permite diferentes níveis de acesso com base em funções de usuário atribuídas aos diferentes níveis de conhecimento especializado na organização. A capacidade de programação simples combinada ao monitoramento de SNMP proporciona mais flexibilidade de gerenciamento. O DDMC oferece uma opção de pré-verificação antes de agendar um upgrade do DDOS para garantir que o ambiente seja compatível com a atualização. Depois que a verificação inicial for concluída, você poderá agendar um upgrade geral, o que permite agendar diversos upgrades do DDOS em vez de fazer atualizações individuais. A configuração de vários equipamentos da série DD é simples com o DDMC porque permite criar e aplicar modelos desse procedimento. Com ataques cibernéticos e ameaças em ascensão, o DDMC pode gerar alertas quando a configuração de um sistema está fora de conformidade. Em caso de erro de upgrade do DDOS, o equipamento voltará automaticamente à versão anterior do sistema operacional, o que diminui o tempo de inatividade do sistema e permite operações contínuas de backup.

Além disso, a série DD conta com um sistema de geração automática de relatórios call-home chamado Autosupport, que envia notificações por e-mail com o status completo do sistema ao Suporte Dell e a uma lista selecionada de administradores. Esse recurso de alerta não invasivo e de coleta de dados permite suporte e serviço proativos sem a intervenção do administrador, simplificando ainda mais o gerenciamento contínuo.

Os equipamentos série DD já estão integrados ao Dell CloudIQ. O CloudIQ oferece uma percepções e análises proativas de desempenho no armazenamento compatível, proteção de dados e produtos hiperconvergentes através de uma interface do usuário.

Complementos de software da série DD

DD Boost

O software DD Boost oferece um nível avançado de integração a aplicativos de backup e utilitários de banco de dados, melhorando o desempenho e a facilidade de uso. A Dell também oferece o plug-in DD Boost File System (BoostFS) com DD Boost para incluir ainda mais suporte a aplicativos, o que dá todos os benefícios do DD Boost para aplicativos que usam NFS na proteção de dados. Em vez de enviar todos os dados ao sistema para processos de deduplicação, o DD Boost permite que o servidor de backup ou o client de aplicativo envie apenas segmentos de dados exclusivos em toda a rede ao sistema.

DD Replicator

O software DD Replicator oferece replicação automatizada, criptografada, baseada em políticas e com uso eficiente de rede para fins de recuperação de desastres e consolidação do backup e do arquivamento em vários locais. O software DD Replicator replica de modo assíncrono apenas dados compactados e deduplicados na WAN. A deduplicação entre locais reduz ainda mais os requisitos de largura de banda quando diversos locais estão replicando para o mesmo sistema de destino. Isso aumenta a eficiência de rede entre todos os locais e reduz os requisitos diários de largura de banda da rede, fazendo com que a replicação baseada em rede seja rápida, confiável e econômica. Para atender a uma ampla variedade de requisitos de recuperação de desastres, o DD Replicator oferece várias topologias de replicação flexíveis, como "full system mirroring" ("espelhamento completo entre sistemas"), "bi-directional" ("bidirecional"), "many-to-one" ("individual"), "one-to-many" ("geral") e "cascaded" ("sequenciado").

Future-Proof Program e Dell Technologies APEX

O Future-Proof Program é um programa dedicado aos clientes que dá mais tranquilidade com garantia de satisfação e proteção do investimento por meio de um conjunto abrangente de programas e recursos tecnológicos de classe mundial para alterações de tecnologia futuras. A série DD faz parte desse Programa Future-Proof. A série DD faz parte do programa APEX da Dell Technologies, o que permite opções de pagamento flexíveis, como os planos Pay as You Go e Pay As You Use e as ofertas de "as a service" inclusas.



Saiba mais sobre a
[série DD](#)



Entre em contato com um
[especialista da Dell Technologies](#)

Dell PowerProtect Data Manager: Microsoft SQL Database Backup and Recovery

April 2023

H18091.9

White Paper

Abstract

This white paper focuses on protecting Microsoft SQL Server using Dell PowerProtect Data Manager, the next-generation data protection platform.

Dell Technologies

Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2019-2023 Dell Inc. or its subsidiaries. All Rights Reserved. Published in the USA April 2023 H18091.9.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Contents

Executive summary	4
Data Manager overview	6
Deployment requirements	9
Security	13
SQL database backup.....	14
SQL database recovery	28
Support for existing SQL agent backups with Data Manager.....	41
Disaster recovery	42
Technical support and resources.....	43

Executive summary

Business case: Challenges Data owners and IT administrators in midsized or enterprise organizations are seeking a scale-out data management software platform to simplify management, capacity growth, deployment, and upgrades. As businesses continue to consume IT resources differently, there is a need for powerful, efficient, and trusted data protection to enable organizations to transform to meet future demands when modernizing their IT environment.

Solution overview Dell PowerProtect Data Manager is defined with built-in deduplication for data protection, replication, and reuse. Data Manager delivers load balancing, enabled with machine learning, to provide optimal deduplication and performance. Data Manager offers efficient data management capabilities across ever-changing IT environments, leveraging the latest evolution of Dell Technologies trusted protection storage architecture.

PowerProtect Data Manager enables Microsoft SQL DBAs to be in control of their backup and recovery practices, using their native tools. Data Manager provides a choice of centralized management to the backup team, and oversight and governance to IT to ensure compliance. Data Manager also keeps data owners nearest to their data, enabling them to protect and manage and restore data as needed from native applications.

Audience This white paper is intended for customer, partners, and prospects who want to understand how Data Manager helps in protecting Microsoft SQL Server.

Revisions

Table 1. Revisions

Date	Part number/revision	Description
November 2021		Initial release
September 2019		Revised for PowerProtect Data Manager version 19.2 release
February 2021		Revised for PowerProtect Data Manager version 19.7 release
May 2021		Revised for PowerProtect Data Manager version 19.8 release
September 2021		Revised for PowerProtect Data Manager version 19.9 release
December 2021		Updated template
May 2022		Revised for PowerProtect Data Manager version 19.10 release
July 2022		Revised for PowerProtect Data Manager version 19.11 release
October 2022	H18091.8	Revised for PowerProtect Data Manager version 19.12 release
April 2023	H18091.9	Revised for PowerProtect Data Manager version 19.13 release

Note: This document may contain language from third-party content that is not under Dell Technologies' control and is not consistent with 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.

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).


Author: Vinod Kumar Kumaresan

Note: For links to other documentation for this topic, see the [Data Protection Info Hub](#).

Data Manager overview

Introduction

Data Manager manages and monitors data protection and replication for Microsoft SQL Server assets through integration with the Microsoft application agent.



Key features for SQL backup and recovery

- Protection for stand-alone SQL Server and SQL Server clustered environments, including Always On availability groups and failover cluster instances
- Ability to restore a single SQL database or multiple databases to the original or an alternate location and restore Always On availability Groups
- Option to set the parallelism for each backup type as a value 1 through 32 in the full, differential, and log fields
- Option to modify the stripe level of a backup at the individual database level
- Improved performance and scaling for virtual machine SQL protection policies
- Centralized restore support for SQL application aware SQL backups and SQL Application Direct backups
- Install MSAPPAGENT agent manually for Windows VM from Data Manager UI
- DD compressed restore option for SQL Application Direct backups
- SQL Application Direct and VM Direct (application aware) support enables performing table-level restores of encrypted databases. ItemPoint 8.6.1 as integrated with the Microsoft application agent now supports the Microsoft SQL Server Transparent Data Encryption (TDE). However, only the AES_128, AES_192, and AES_256 encryption algorithms are supported.
- Dynamic protection and AAG selection options during SQL policy creation in Data Manager UI
- Option to exclude simple recovery model database and system database in protection policy

- Notification for the assets that are skipped from backup when the assets are offline or in a restoring, recovery pending, or suspect state
- Option to configure Windows firewall during the application agent installation (open port 7000 by adding firewall rule during agent installation)
- Dynamic protection rule to ensure that the protection policy dynamically protects all the selected assets within the selected host or instance container
- Exclude unprotectable database - An advanced option in the Data Manager UI Add Policy wizard to skip the SQL databases that cannot be backed up
- Option to set up the sysadmin privilege for a database OS account for a host with application agent release 19.11 or later installed
- Improved Microsoft SQL Server backup promotion logic when transaction log backup chains span across full backups - When Microsoft SQL Server transaction log backup chains span across full backups, the transaction log backups and full backups can run concurrently. The transaction log backups are not promoted to full backups.
- Starting with version 19.12, PowerProtect Data Manager supports SQL protection on virtual machines with UAC enabled
- Table-level recovery of Microsoft SQL Server backups with ItemPoint on Windows Server 2022 using Microsoft application agent version 19.12.
- Microsoft SQL Server 2022 support – Starting with version 19.13, PowerProtect Data Manager supports both centralized and self-service Application Direct operations to protect Microsoft SQL Server 2022 databases. PowerProtect Data Manager supports the backups and restores of Microsoft SQL Server 2022 databases through the Microsoft application agent, including table-level restores that use ItemPoint.
- Streamlined restore of Always On Availability (AAG) groups in PowerProtect Data Manager UI - The AAG database restore automatically removes the database from the AAG, restores the database to either a single server instance or all the nodes in the AAG, as specified in the UI, and then adds the database back to the AAG.
- Enhanced restore wizard - PowerProtect Data Manager UI screens are enhanced for the Microsoft SQL Server restore workflows, to improve restore wizard usability. The PowerProtect Data Manager UI includes new options for restore to an original or alternate database. It also provides both hierarchical and list views for the restore to an alternate database, where the hierarchical view displays the databases for each instance on the available Microsoft SQL Server hosts.

Architecture

Before we dive into Data Manager for Microsoft SQL Database backup and recovery, let us understand the key architectural components involved in SQL data protection.

SQL Server Management Studio (SSMS) plug-in

The Microsoft application agent for Application Direct with SQL Server has an SQL Server Management Studio (SSMS) plug-in. The plug-in is similar to the SQL native backup and restore UI.

Virtual Device Interface (VDI)

The Microsoft Application Agent for Application Direct with SQL Server uses a VDI (an API provided by SQL Server) to integrate with the SQL Server. It also enables the Application Direct with Microsoft application agent to back up and restore SQL Server data.

DD Boost library

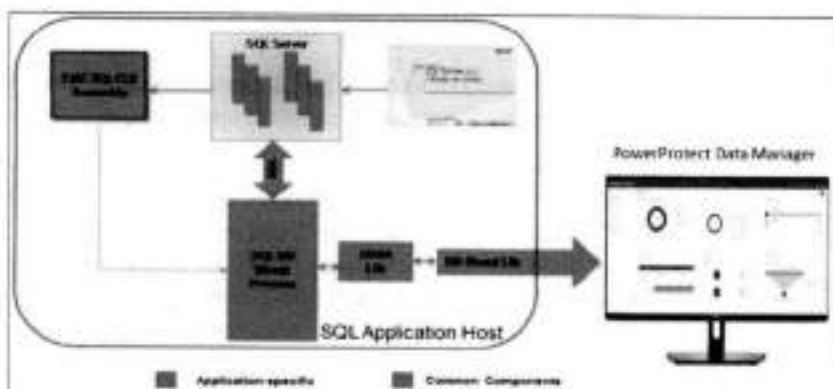
The DD Boost library performs source-based deduplication and sends the backup data to Dell PowerProtect DD series appliances.

Microsoft application agent

The Microsoft application agent receives requests from the vProxy engine (using vProxy agent) to perform a certain set of operations during a backup to:

- Achieve application consistency when taking snapshots (full image)
- Perform transaction log backup

The application agent enables database administrators to back up and restore SQL databases using SSMS.



SQL-CLR assembly

CLR and T-SQL scripts must be integrated to create functions or procedures to perform backups and restores in a SQL environment. The Microsoft application agent installation deploys the CLR assembly. The CLR assembly contains one exportable SQL function type routine to run any Microsoft application agent command at the command prompt.

Lockbox

The lockbox is an encrypted file that the Microsoft application agent uses to store confidential data, such as login credentials, and to protect that data from unauthorized access.

Registering a DD series appliance to a new lockbox creates a lockbox folder. Registering a server to or removing a server from the lockbox updates the PersistedSettings.xml file. The PersistedSettings.xml file contains the DD series appliance information, such as the server name, communication protocol, FC service name, username of the DD Boost user, and storage unit.

Lockboxes were configured by default in the path C:\Program Files\DPSAPPS\common\lockbox until 19.8 version release. Starting from Data Manager

v19.9 release, the lockbox is automatically configured in the custom path where the agent is installed.

Recommendations for lockbox configuration

For a DD series appliance, you can configure one of the following lockbox types according to the environmental requirements:

- **Single lockbox:** In a stand-alone environment, create a single lockbox on the host.
- **Shared lockbox:** In an environment with multiple instances, configure a single lockbox in a shared location and grant each remote host individual access to the lockbox.

Note: Do not use a single shared lockbox to perform remote backup operations in a cluster environment. The backup will fail. Use multiple lockboxes in a cluster environment.

- **Multiple lockboxes:** In an environment with multiple instances, configure a lockbox on each instance in the environment.

vProxy engine

A vProxy engine is a data mover that takes a snapshot of VM and copies the data from VMware datastores to DD series appliance using vSphere APIs for Data Protection (VADP). A vProxy protection engine uses the VMware snapshot technology and uses vSphere Web Service APIs to take the VM snapshot and leverage off Changed Block Tracking (CBT) during the backup.

vProxy agent

A vProxy engine deploys its own vProxy agent on a guest VM. The vProxy engine uses this agent through a VMware infrastructure communication protocol to install and talk to the Microsoft application agent in the guest VM to perform app-aware backups. This does not require any IP network connectivity because it goes through VMware infrastructure using VMware tools.

Deployment requirements

Network

To protect Microsoft SQL Server using Data Manager, follow these key network requirements:

- All clocks on both the SQL Server host, domain controller, and Data Manager server are time-synced to the local NTP server to ensure discovery of the backups.
- The SQL Server and Data Manager network can see and resolve each other.
- Port 7000 is open on the SQL Server host and is bi-directional.
- Port 8443 is open on the SQL Server host. This requirement applies only to VM Direct.
- DNS is configured correctly on the application agent host for SQL Server.
- DNS is configured correctly on the Data Manager host, and the name resolution matches.

Microsoft application agent The Microsoft application agent is a backup and recovery agent that is built for Microsoft applications.

The Microsoft application agent integrates with the Microsoft SQL Virtual Device Interface (VDI) to enable application owners to protect the Microsoft SQL database. The Microsoft application agent provides a user interface (Microsoft application agent for Application Direct) that enables the Microsoft SQL database administrator to perform backup and restore operation of the SQL database through the SSMS Plug-in. The plug-in is like the SQL native backup and restore user interface. The Microsoft SQL database administrator can use the Microsoft application agent for Application Direct to back up and restore the Microsoft SQL database.

With Data Manager, the Microsoft application agent supports multiple tools to perform manual backups. The Microsoft SQL database administrator can configure and perform SQL databases backup using the following tools:

- Microsoft application Agent for SSMS Plug-in (Microsoft application agent for Application Direct)
- Microsoft application Agent for Application Direct commands
- T-SQL scripts

SQL host

On the Microsoft SQL Server host, install the Microsoft application agent and configure the installation options. Download the Microsoft application agent .zip file, from the [Dell Technologies support site](#) to the Microsoft SQL Server host. Extract the .zip file and run the .exe file to start the installation wizard.

Prerequisites for Always on Availability Group (AAG) protection

Agent installation: For Microsoft SQL Always on Availability environment, the Microsoft application agent needs to be installed on each node in the cluster. If a database is protected in an Always on Availability Group (AAG), stand-alone backups cannot be configured of that database in a protection policy group. SQL AAG requires a Microsoft Cluster Service (MSCS) starting from SQL Server 2017.

Lockbox configuration: Either a single shared lockbox or a separate lockbox must be configured on each node that is in the Always on Availability Group.

Set Readable Secondary settings: Set the Always on Availability Group readable secondary configuration option to either yes or read intent only:

1. On SSMS, in the Object Explorer, right-click the Always on Availability Group and select Properties. The Availability Group Properties dialog box appears.
2. In the availability replicas table, readable secondary column, select either yes or read intent only for each of the primary and secondary replicas of the SQL Server instances. This setting allows the Microsoft application agent to gather information about the secondary replica (for example, database file location, which can be different from the other replicas).

See [PowerProtect Data Manager Compatibility Guide](#) for supported Microsoft application agent versions and Microsoft SQL Server versions with the Data Manager.

Data Manager support for SQL AAG and SQL cluster-less AAG environment

Data Manager supports both SQL AAG and SQL cluster-less AAG environment for database protection and recovery.

SQL AAG	SQL cluster-less AAG
<ul style="list-style-type: none"> ✓ Provides both high availability (HA) and Disaster Recovery (DR) ✓ Can host one primary and up to 8 secondary replicas ✓ Primary replica hosts a read write copy while the secondary copies can be read only ✓ Zero data loss protection using synchronous data replication ✓ Automatic or manual failover ✓ Local or shared storage ✓ Automatic page repair ✓ Active use of secondaries 	<ul style="list-style-type: none"> ✓ Not an HA solution ✓ Can host one primary and any number of secondary replicas ✓ Primary replica hosts a read write copy while the secondary copies can be read only ✓ No automatic failover, only manual failover without data loss and forced failover with data loss is possible ✓ Supports only read scale workloads ✓ Can include replicas that are hosted on a variety of operating system platforms

Database discovery on Data Manager

Database discovery	
Database backup	Value
Standalone database backups	Local hostname
Clustered Availability group database backup	Cluster name
Cluster less Availability group database backup	AGName_AGGUID'
SQL virtual server database backup	Network name

Clustered SQL AAG database discovery on Data Manager

The following figure represents the SQL AAG database which is shown as assets for data protection.



Cluster-less SQL AAG database discovery on Data Manager

The following figure shows that the cluster-less SQL AAG database appears as assets for data protection.



The following features are supported with Application Direct and VM application-consistent backups:

- Centralized backups from Data Manager
- Self-service backups (only Application Direct)
- Self-service database restore options
 - Point in Time (PIT) restore
 - Restore to latest point

- Table-level restore
- Instance access of the database (only VM Direct)
- Alternate restore
- Flat-file restore

Security

SQL authentication

The Microsoft application agent requires that the user starting backup and recovery operations are assigned certain privileges from the SQL Server and the Windows application host. The following table explains the required permissions for both stand-alone SQL and AAG SQL Servers.

SQL Server	Required SQL Server roles	Required Windows user permissions
Stand-alone SQL Server	sysadmin and public	<p>Create a local or domain windows user account and assign the following roles:</p> <p>For table-level backup and recovery, assign administrative privileges.</p> <p>For database-level backup and recovery, assign administrative permissions:</p> <ul style="list-style-type: none"> • Add the user to the 'create global objects' windows policy. • Assign the following permissions to the data and log folder of the database: <ul style="list-style-type: none"> ▪ Read ▪ Write ▪ List folder contents
Always-on Availability Group	sysadmin and public	<p>Create a Windows user account with one of the following configurations</p> <ul style="list-style-type: none"> • Domain user added to the administrator's user group. • The built-in windows administrator. • The local user account added to the administrator's user group of each node in the cluster. The username and password must be the same on each node.

The Microsoft application agent supports SQL data encryption at the cell level, at the full database level by using TDE, or at the file-level with encryption options provided by Microsoft. (Microsoft SQL transparent data encryption (TDE) is a feature that performs real-time I/O encryption and decryption of the data and log files.)

Note: The Microsoft application agent does not support third-party transparent data encryption for SQL VDI. See the Microsoft SQL Server product documentation for more information about TDE, enabling data encryption, and protecting the encryption keys.

Sysadmin privilege for Microsoft SQL Server hosts

To enable the integration with PowerProtect Data Manager on each Microsoft SQL Server host, ensure that the database OS account or NT AUTHORITY\SYSTEM account on each host has the required sysadmin privilege.

Note: Setting the sysadmin privilege for the database OS account is only supported for a host with application agent release 19.11 or later installed.

For more information about how to set the sysadmin privilege for database OS or NT AUTHORITY\SYSTEM account, see the section "Setting the sysadmin privilege for Microsoft SQL Server hosts" in the [PowerProtect Data Manager Microsoft SQL Server User Guide](#).

SQL protection on virtual machines with UAC enabled

With VM Tools version 11.x or later installed on the virtual machine, you can configure the required protection policies and perform the VM Direct backups and restores by using domain user or local user credentials with enabled UAC. For more information, see [PowerProtect Data Manager SQL Server User Guide](#).

SQL database backup

Data paths

Data Manager provides flexibility to protect the Microsoft SQL database using the methods below. Data Manager supports the following data paths for Microsoft SQL database protection:

- Centralized protection using Application Direct data path
- Self-service protection using Application Direct data path
- Centralized virtual machine application-aware using VM proxy data path

Centralized protection policy: When the admin creates a protection policy for Microsoft SQL database, the centralized protection option enables Data Manager to centrally manage the entire life cycle of data protection operations for Microsoft SQL database.

Self-service protection policy: When the admin creates a protection policy for Microsoft SQL database, the self-service protection option enables the data owner to perform the manual backup operation from the command line interface.

Data Manager prepares the environment to accommodate the manual backup operation, such as:

- Creating a user with password for data protection storage
- Creating a storage unit
- Enforcing the backup data retention

The following data protection attributes are specified when the self-service protection policy is created: Application Type, Purpose, Assets, Schedule, and SLA

Note: Only the retention period can be specified in the schedule attribute in the self-service protection policy.

Centralized virtual machine application-aware using VM proxy data path: If the Microsoft SQL Server is running in a ESXi virtual machine, the application-aware option enables Data Manager to interact with the Microsoft SQL Server virtual machine and install the Microsoft application agent for an application-consistent backup of Microsoft SQL database.

The Microsoft application agent is a component of the vProxy data protection solution that is bundled with the vProxy appliance. The vProxy automatically deploys the agent during a virtual machine application-aware backup and, if required, when restoring Microsoft SQL databases and SQL instance backups to running virtual machines.

After installation, the Microsoft application agent package appears in the Windows installer Add-Remove programs list.

The Microsoft application agent allows for advanced application data protection of workloads residing on a VMware ESXi server. This includes adding SQL virtual machines to an advanced application-consistent protection policy to perform the following operations:

- **SQL Server FULL backup:** Configure protection policy with the application-aware option to perform SQL Server backup to DD series appliance as part of a VMware image-level backup. The SQL Server FULL backup is performed during the in-guest quiesce by VMware Tools. After running the policy, the catalog and index information for the SQL server backup is stored on the DD series appliance. When the backup is performed as part of the VMware image-level backup, the SQL data files are backed up as part of the VMDKs during the vProxy image backup.
- **Transaction log backup:** When configuring protection policy with the application-aware option, set an interval for transaction log backup to enable transaction log backups for SQL instances running in the virtual machine and specify the frequency of backups. Backups are written directly to DD series appliance. Transaction log backup is only performed for databases in the proper state, otherwise databases are skipped.

Note: The application-aware backup option requires vSphere version 6.5 or 6.7 and VMware tools version 10.1 or later. Selecting the application-aware type of backup Microsoft SQL virtual machine also enables the recovery of Microsoft SQL databases by using the SSMS.

SQL database backup



To add an application-aware protection policy, select the Virtual Machine type, and then select Application Aware and provide the Microsoft SQL Server login credentials as shown above.

Native view for SQL assets

Native asset view uses a SQL tree view that shows the hierarchical relationships of the SQL hosts, their application servers, or instances (including Failover Cluster Instances (FCIs)), stand-alone database assets, and any Always On availability groups (AAGs) with their database assets. To list all the assets and AAGs within a host and instance, expand the hierarchical or tree view. Selecting a host or instance container also selects all the contained assets and objects. You can also select either individual assets or a group of assets within the host or instance container to include in the protection policy.



Note: The native view is supported only for an Application Direct protection policy and not for an application-aware protection policy.

Dynamic protection

Enabling this icon (enabled by default), automatically creates a dynamic protection rule to ensure that all the selected assets within the selected host or instance container are dynamically protected by the protection policy. Data Manager manages the protection rule. Any of the following actions updates the rule: editing the policy, making changes to the container selections, or moving assets into or out of a selected container.

After creating or editing the policy, dynamic protection ensures that any new databases, or assets that are added to the instance or AAG or container, are added automatically to the protection policy. Removing any databases or assets from the instance or AAG or container also removes them from the policy. When any selection overlap occurs between different policies, the UI displays the overlaps and helps to resolve the SQL asset assignment conflicts by adjusting the protection rules' priority.



To disable the dynamic protection for a container, click the Dynamic Protection icon and then click Disable in the displayed text box. When the dynamic protection is disabled, the protection policy does not dynamically protect the selected container and its objects. As a result, all the selected objects within the container become static selections that the policy does not automatically protect.

After creating or editing a protection policy that has dynamic protection, go to **Protection > Protection Rules** to see the protection rule details for the protection policy, including the priority of the protection rule. Dynamic protection rules apply only at the container level.

Excluding databases from backup

Starting with Data Manager version 19.9, protection policy enables the option to exclude databases from backup.

Job ID	Status	Description	Policy Name
2N7C8W	Completed with Exceptions	Manual Protecting SQL Databases - SQL APP Direct - PROTECTION - Differential	SQL APP Direct

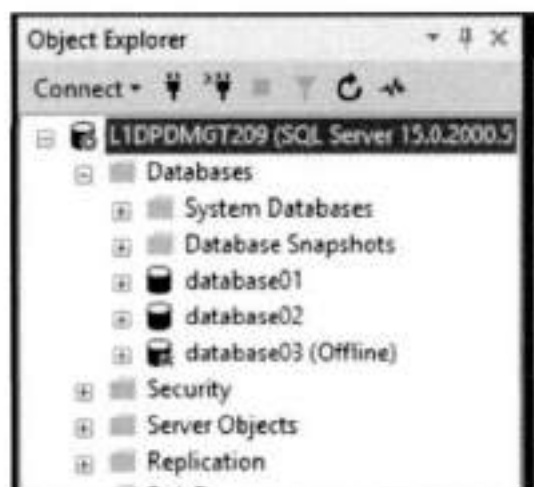
The reason for the **Skipped** status appears in the details section of the **Job ID Summary** window:



Exclude unprotectable database

Until Data Manager version 19.9, if a database was in an unprotectable state (offline / recovery mode), its backup would fail. Starting with Data Manager version 19.10, you can exclude the unprotectable SQL databases from all backups. This is a protection policy option to exclude unprotectable SQL databases from the Data Manager UI.

For example, the following SQL Server has two online databases (database01 and database02) and one offline database (database03).



Select **"Exclude Unprotectable Database"** on the **Protection > Protection Policies > Add Policy > Options** page.

SQL database backup



By selecting the option **"Exclude Unprotectable Database"** in the protection policy, Data Manager provides the ability to skip the unprotectable database (offline / recovery mode).

Once the backup is complete, the reason for the **Skipped** status appears in the errors section, as shown here:



Asset assignment conflict notification

If the **Assets** page included any asset assignment conflicts with other protection policies, a notification page appears that describes the assets that are already assigned to other protection policies.

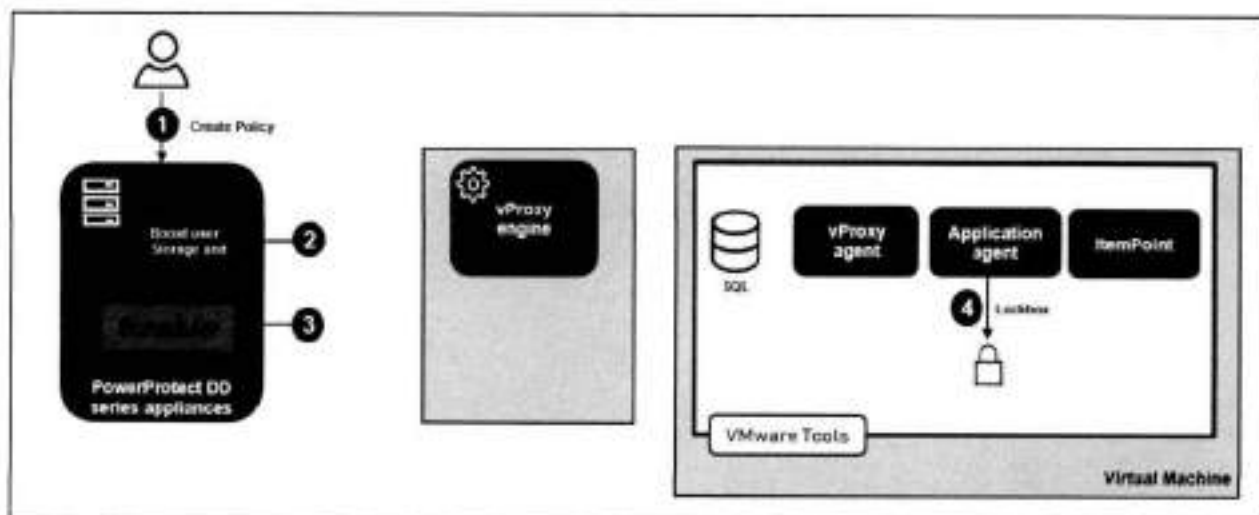


Selecting **Next** displays the **“Check conflicts due to rule priority”** page that shows the assets that have conflicting assignments and their protection policies and rules. In the **Protection Rules** pane, use the up and down arrows to change the protection rule priority of any policy. Raising the rule priority for a policy moves the assets with conflicts in a lower-priority policy to the policy that has the higher protection rule priority.



Centralized Application Direct backup workflow

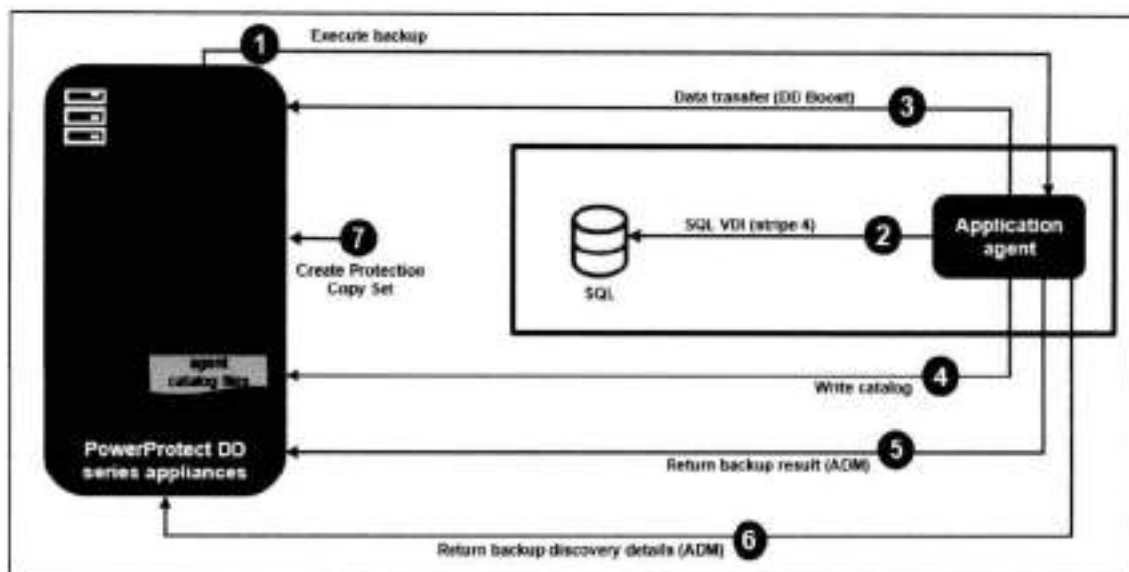
Configuration stage - The configuration stage of central Application Direct backup consists of the following steps.



1. The user creates a protection policy from Data Manager UI.
2. Data Manager creates a boost user and storage-unit on DD series appliance.
3. Data Manager adds a protection schedule to its own scheduler.
4. The application agent configures a lockbox with credentials on the SQL host.

Centralized Application Direct backup workflow (FULL)

Protect stage – The following figure provides details about the centralized Application Direct backup flow.




1. Request to the App agent to execute a backup.
2. The app agent uses SQL VDI (stripe 4) to back up SQL.
3. The app agent moves data to DD series appliance using DD Boost.
4. The Microsoft app agent catalogs the backup.

(Repeat steps 2 through 4 for each database.)

5. Return backup results to PowerProtect (using ADM).
6. Return backup discovery details to PowerProtect (using ADM).
7. PowerProtect creates SQL PCS (Protection Copy Set) based on the backup results.

Centralized Application Direct backup workflow (LOG)

Protect stage – The following figure provides details about the centralized Application Direct backup flow.



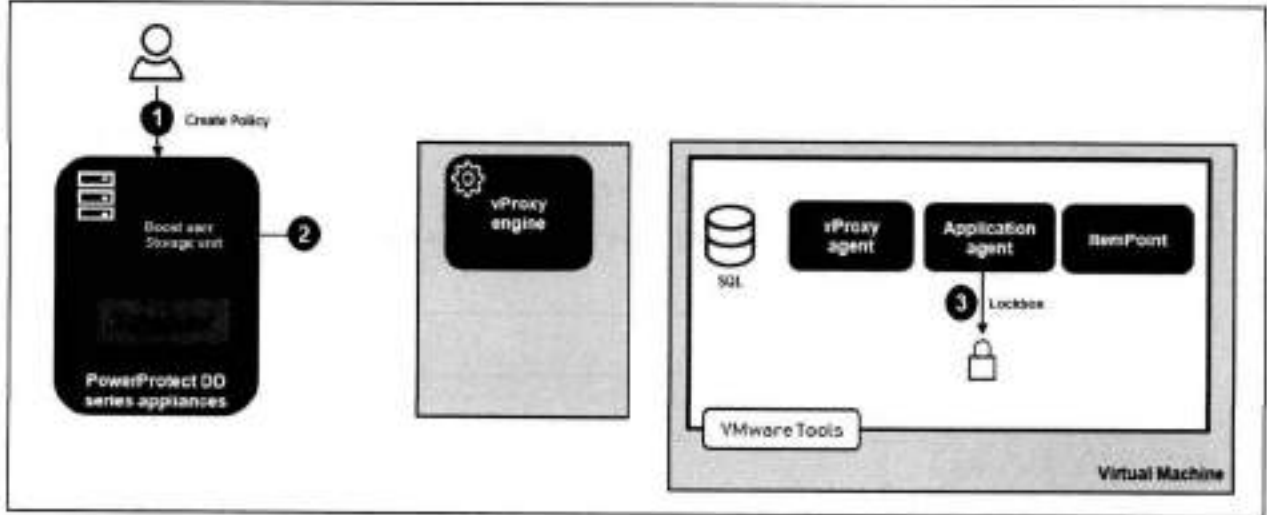
1. Request to the App agent to execute a backup.
2. The app agent uses SQL VDI (stripe 1) to back up SQL.
3. The app agent moves data to DD series appliance using DD Boost.
4. The Microsoft app agent catalogs the backup.
5. The SQL Server truncates the log.
(Repeat steps 2 through 5 for each database.)
6. Return backup results to PowerProtect (using ADM).
7. Return backup discovery details to PowerProtect (ADM).
8. PowerProtect creates SQL PCS (Protection Copy Set) based on the backup results.

Self-service Application Direct backup workflow

Configuration stage - The configuration stage of self-service Application Direct backup consists of the following steps:

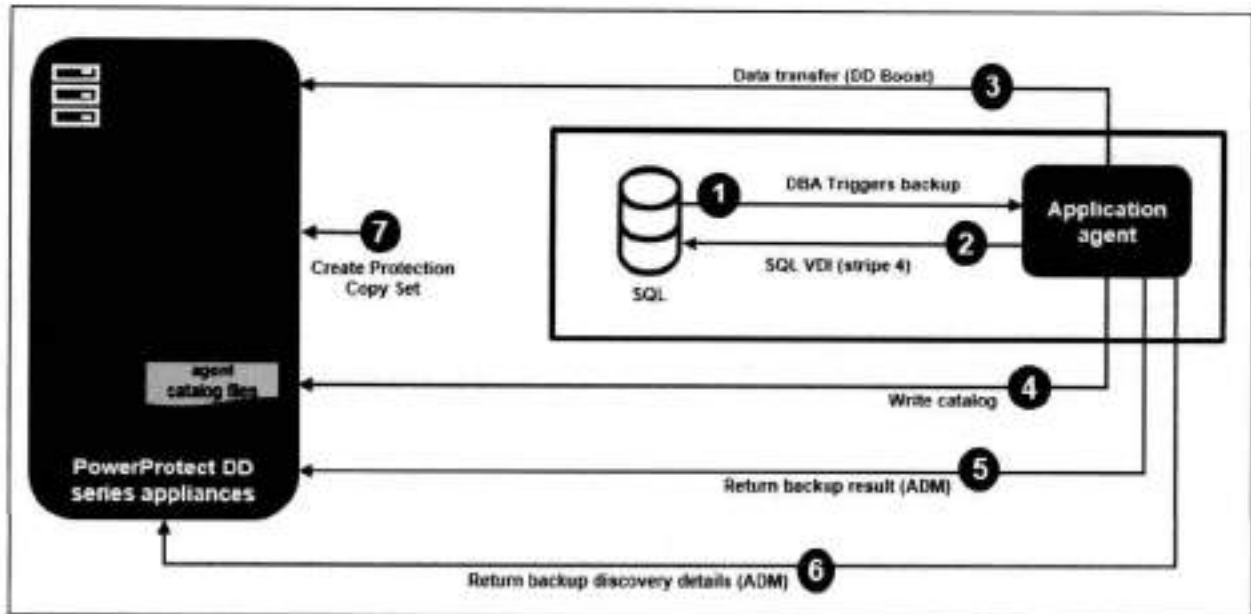
1. The user creates a protection policy from Data Manager.
2. Data Manager creates a Boost user and storage-unit on the storage.

3. The application agent configures a lockbox with credentials on SQL.



Self-service Application Direct backup workflow (FULL)

Protect stage - The following steps show the self-service Application Direct SQL database backup flow.

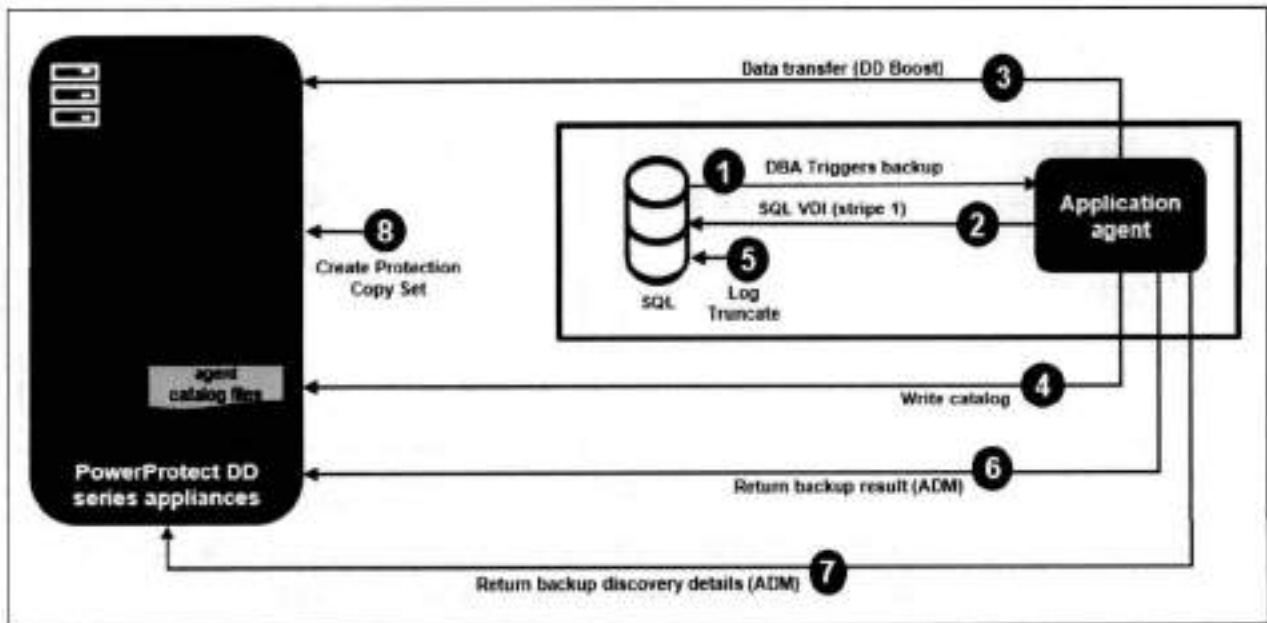


1. Protection triggered by DBA on SQL host.
2. The app agent uses SQL VDI to back up SQL.
3. The app agent moves data to DD series appliance using DD Boost.
4. The Microsoft app agent catalogs the backup.
5. Return backup results through ADM.
6. Return backup discovery details (ADM).

- PowerProtect creates SQL PCS (Protection Copy Set) based on the backup results.

Self-service Application Direct backup workflow (LOG)

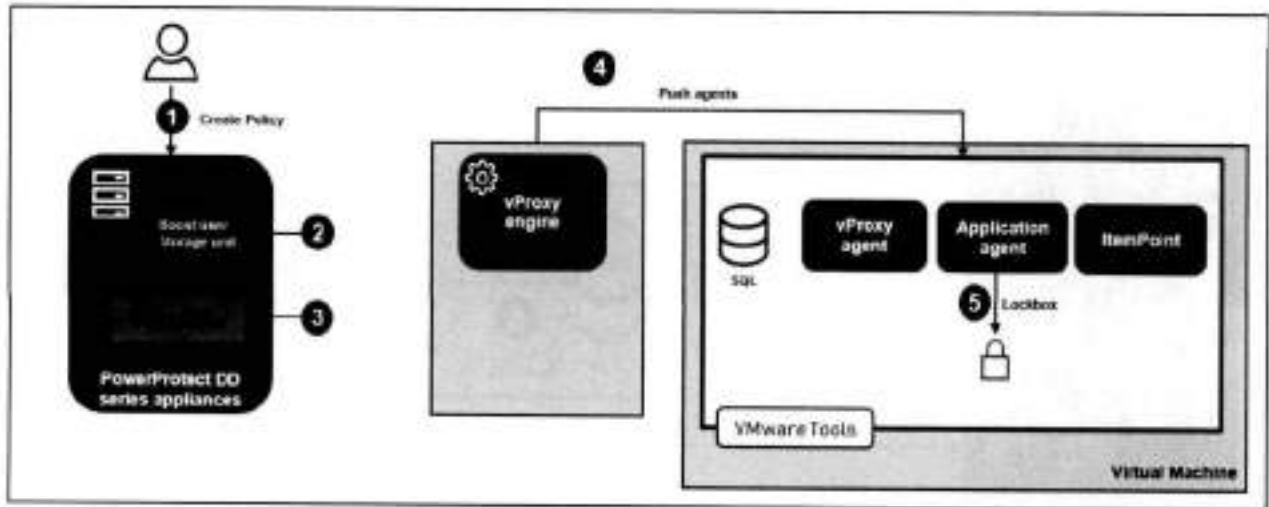
Protect stage - The steps below show the self-service Application Direct SQL log backup flow.



- Protection triggered by DBA on SQL host.
- The app agent uses SQL VDI to back up SQL.
- The app agent moves data to DD series appliance using DD Boost.
- The Microsoft app agent catalogs the backup.
- The SQL Server truncates the log.
(Repeat steps 2 through 5 for each database.)
- Return backup results through ADM.
- Return backup discovery details (ADM).
- PowerProtect creates SQL PCS (Protection Copy Set) based on the backup results.

Application-aware backup workflow

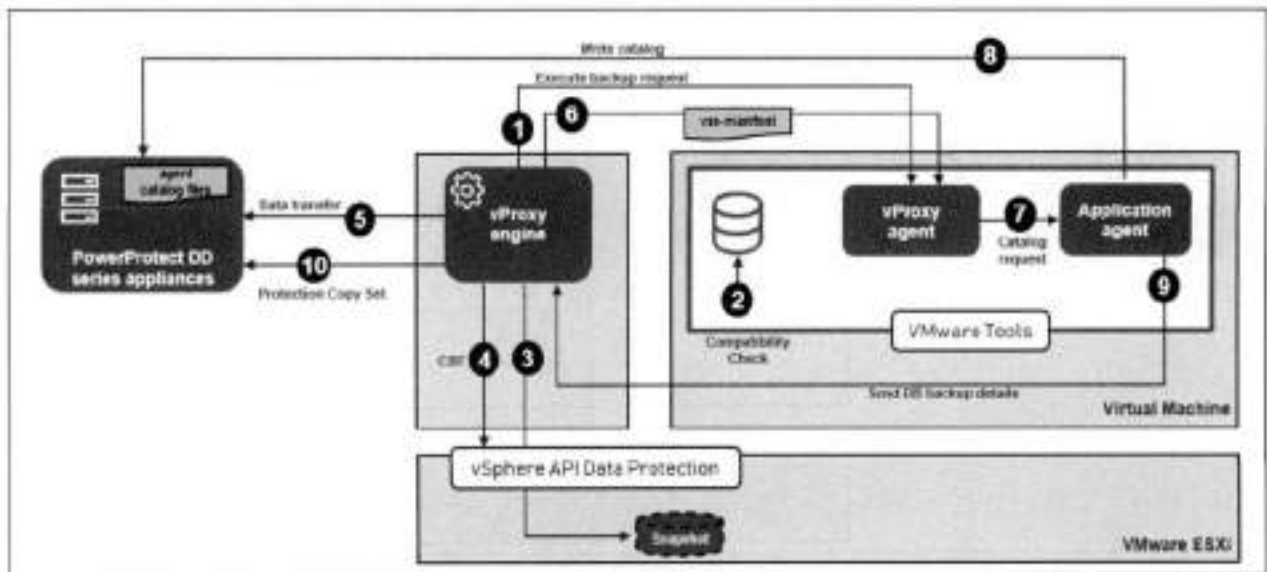
Configuration stage - The configuration stage of application-aware backup consists of the following steps.



1. The user creates a protection policy using PowerProtect UI.
2. PowerProtect creates a Boost user and a storage-unit on DD series appliance.
3. PowerProtect adds a protection schedule to its own scheduler.
4. The vProxy engine pushes agents using guest operating system credentials:
 - Microsoft application agent
 - Dell vProxy agent
 - Dell ItemPoint
5. The application agent configures lockbox with credentials on the SQL host (using ADM).

Application-aware SQL Database Backup Workflow (FULL)

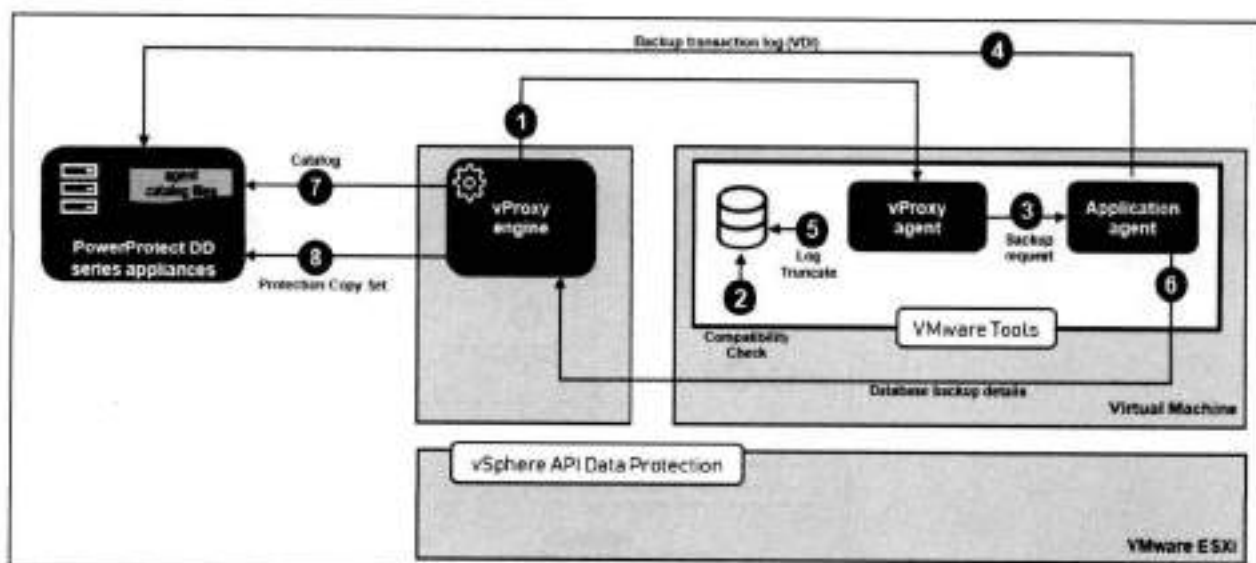
Protect stage - The following figure shows the steps involved in the application-aware backup for database.



1. Request to the vProxy agent to execute the backup.
2. Application-aware compatibility check (SQL permission, SQL status, VSS status, and so on).
3. The vProxy takes VADP snapshot with quiesce option which will internally trigger VMware's own VSS workflow.
4. The vProxy gets changed blocks from VADP.
5. The vProxy starts the data transfer to DD series appliance.
6. The vProxy retrieves the VSS manifest (metadata) from vSphere using VADP API and uploads it to the guest VM.
7. The vProxy tells the Microsoft app agent to catalog the backup.
8. The app agent parses the VSS manifest and catalogs databases quiesced during step 3 under its own directory structure.
9. The app agent provides database backup details, including discovered SQL assets, to the vProxy.
10. PowerProtect creates VM PCS (Protection Copy Set) and a corresponding SQL PCS based on the backup results.

Application-aware backup workflow (LOG)

Protect stage - The following figure shows the steps involved in the application-aware backup of SQL log protection.



1. Request to vProxy to execute the backup.
2. Application-aware compatibility check (SQL permission, SQL status, VSS status, and so on).
3. The vProxy asks the Microsoft App agent to execute the transaction log backup.
4. The Microsoft App agent will serially back up each database transaction log (using VDI) on DD series appliance.
5. The SQL Server truncates the logs.
6. The app agent provides database backup details, including discovered SQL assets, to the vProxy.
7. The vProxy parses the VSS manifest and catalogs files and transaction logs.
8. PowerProtect creates VM PCS and its corresponding SQL PCS based on the backup results.

SQL database recovery

Database recovery overview

The Microsoft application agent provides a user interface (Microsoft application agent for Application Direct) that enables the Microsoft SQL database administrator to perform backup and restore operation of an SQL database through the SSMS plug-in.

For database or table-level restores directly from a SQL host using the SSMS plug-in with Data Manager, the Microsoft application agent supports multiple tools to recover Microsoft SQL databases.

The Microsoft SQL database administrator can use the following tools to configure and perform a SQL database recovery:

- Microsoft application agent SSMS plug-in (Microsoft application agent for Application Direct)

- Microsoft application agent for Application Direct commands
- T-SQL scripts

For table-level restore (through ItemPoint)

- Application agent will mount the image backup
- Restore the required transaction logs
- Replay logs to the selected point-in-time

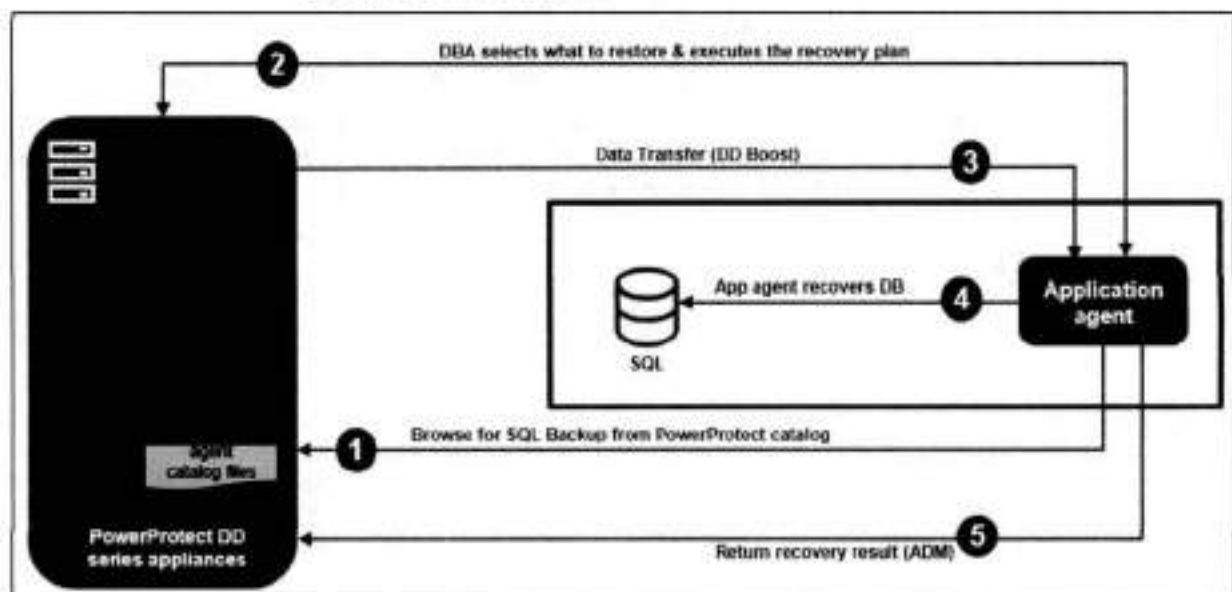
Restore workflow

1. Restore databases using the SSMS plug-in.
2. Browse SQL backups using the application agent catalog stored on Data Manager.
3. DBA selects what to restore and recover.

The application agent creates and executes the recovery plan.

The application agent also restores the transaction logs using DD Boost and applies the transaction logs on the restored database files accordingly.

4. The application agent recovers and opens the database unless the administrator selected a different option.



Database or transaction logs restore as a flat file

- Does not restore data to the active SQL instance
- Files are restored as a flat file without impacting any active SQL server
- DBA will restore and recover data, using these flat files and SQL Server commands

SQL instant access

An instant access recovery enables you to quickly bring a database online from a point-in-time by running the database directly on protection storage, which is added to the host through a backup image mount.

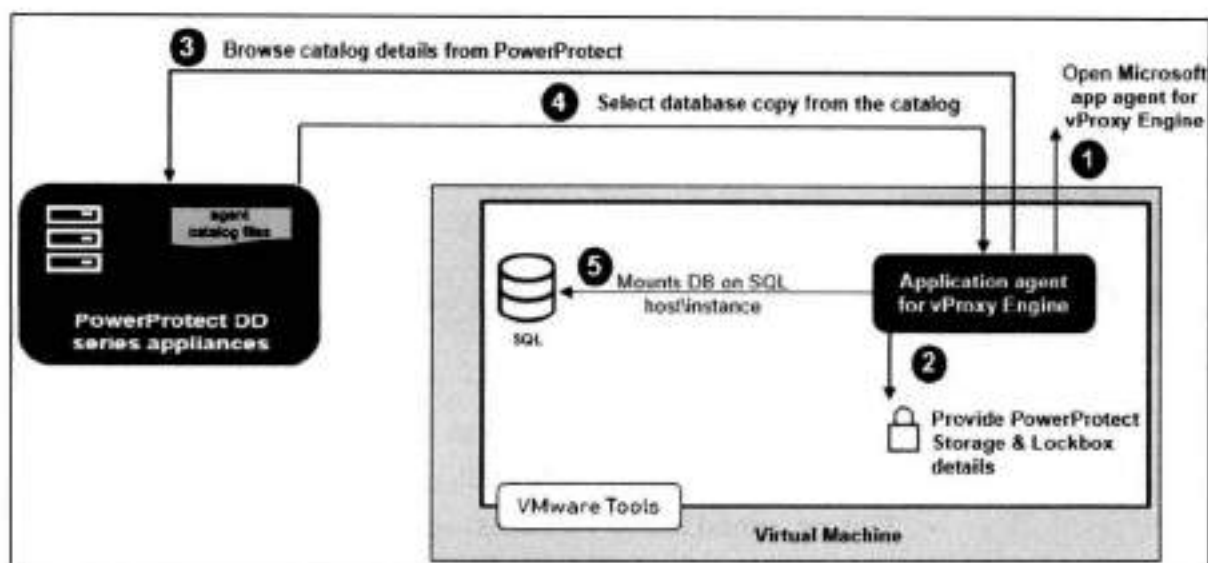
Instant access recovery enables you to access a live-mounted SQL database from protection storage without restoring the virtual machine or SQL database. This type of operation is useful for database administrators to query a SQL database before restoring the database due to time and resource constraints. This feature also supports the migration of mounted disks containing live mounted databases from protection storage to production datastore.

Important: Instant access is only supported for application-aware copy, and not for Application Direct backups.

When you initiate instant access recovery, the operation locates the corresponding backup virtual disks and mounts them from the DD protection storage. The operation locates the SQL Server database selected by the database administrator from mounted disks and connects the database to the SQL Server instance. If transaction logs were selected, the transaction logs are replayed against the instant access database.

The following figure shows the steps for SQL Instant access using the vProxy engine SSMS plug-in:

1. SQL Instance access can be initiated from Microsoft application agent for vProxy engine.
2. From the Instant access page, provide the DD series appliance, lockbox settings, and SQL Instance server.
3. Browse for the available SQL backups in the catalog and select backup date.
4. When a single database is selected, you can specify a new name for the database. For instant access, a default name is generated, by appending the text "Livemount" and a date/time stamp to the original database name.



Centralized restore of SQL Application Direct backups

When the SQL Server data is backed up as part of a SQL Application Direct protection policy, the SQL Server Application Direct backups can be recovered using the centralized restore functionality in the Data Manager UI.

Types of centralized restores of SQL Application Direct backups, depending on the type of database assets:

- Centralized restore of a system database
- Centralized restore of a stand-alone database
- Centralized restore of an Always On availability group (AAG) database
- Centralized restore of a Failover Cluster Instance (FCI) database
- Centralized restore of an Always On Failover Cluster Instance (AAG over FCI) database

Note: The centralized restore of a SQL Application Direct backup and the centralized restore of a SQL virtual machine backup cannot be performed simultaneously. For Centralized restore of an Always On availability group (AAG) database and Failover Cluster Instance (FCI) database, the database needs to be taken out of the cluster.

You can restore single or multiple databases from the same SQL host and instance. The databases can be restored either to the original SQL host or to an alternate SQL host with the following requirements:

- The alternate host must be a SQL Application Direct machine.
- The Microsoft application agent software must be installed and configured on the alternate host.
- A system database cannot be restored to an alternate host or SQL instance.

A centralized restore of a full, differential, or transaction log backup can be performed to a specified SQL host and instance. Centralized restore of SQL Application Direct backups can be performed from the **Recovery > Assets > SQL** window in the Data Manager UI.



Restore to original database:

Databases can be restored to the original database on the original SQL Server when the SQL host is protected by SQL Application Direct policy and SQL host is running and operational. Backup copies are used for operational restore and for Disaster Recovery purposes.

For a single database restore, the restore can be performed from a most recent copy, a specific copy, or from a point-in-time (PIT) copy.

Copy and point-in-time selection:

A single database restore supports copy and point-in-time (PIT) selection. With copy selection, FULL, DIFF or TLOG can be selected and with PIT selection, roll-forward time can be selected and available only when TLOG backup is selected.



Note: Select roll-forward time in between the timelines of full backup copy and log copy.

Restore to an alternate database:

Databases can be restored to an alternate SQL server where the SQL host may (or may not) be protected by Data Manager and the Application agent must be installed manually and registered for "unprotected" host.

Data Manager provides a rich set of options for restoring SQL databases to alternate locations. For an alternate restore, you can:

- Select the target host
- Select the target SQL Instance
- Select a target database



Refresh option to discover recently added SQL instances and databases on the host.

The target host must have the application agent installed and registered by the host administrator. Data Manager will automatically discover the target hosts, SQL instances, and SQL databases.

Database file location:

Use the following restore options to specify the file system location where the databases are restored:

- **Restore database files to the original file location (location at backup time)** - Restores the backup data to the file directory that was used during the backup and overwrites the existing contents.

Note: When restoring to the original path, the file system paths must exist and filenames must not be in use.

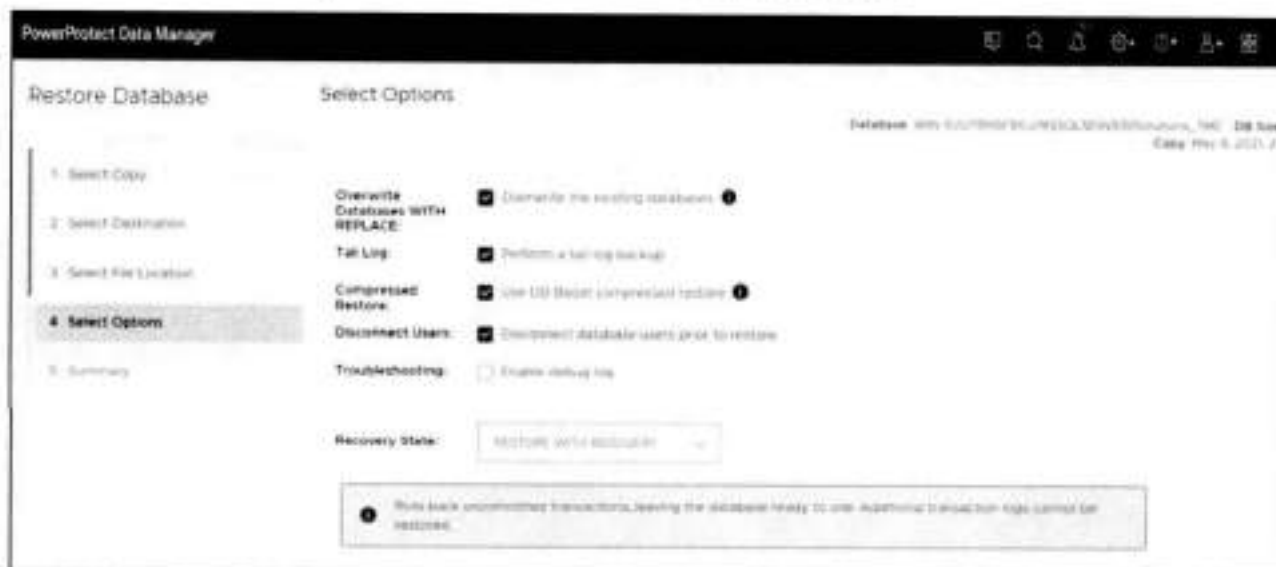
- **Restore database files to the default file location as set by Microsoft SQL Server** - Restores the backup data to the default file directory as used by the SQL Server.
- **Restore database files to a user-specified file location** - Restores the backup data to a user-specified file directory for the database files and log files.

Note: When restoring to a custom location: a) the user can specify one path for data files, and one path for log files, and b) the file system paths must exist.

Tail log backup:

Tail log is supported only for a single database, for backing up the active portion of the database log present on the disk prior to starting the database recovery.

There is an option to enable DD compressed restore, overwrite existing databases, TLOG backup and to disconnect database users prior to restore.



Compressed restore option - Users can use the compressed restore option for saving network bandwidth. However, the compute resource utilization would be more for compressing/decompressing the data.

For restoring multiple databases, a restore can be performed only from the most recent backup copy.

See [PowerProtect Data Manager Microsoft Application Agent SQL Server User Guide](#) for more details.

Centralized restore of SQL virtual machine backups

When the SQL Server data is backed up as part of a SQL virtual machine application-aware protection policy in Data Manager, the SQL server virtual machine backups can be restored using the centralized restore functionality in the Data Manager UI.

Types of centralized restores of SQL virtual machine backups:

- Centralized restore of a system database
- Centralized restore of a stand-alone database
- Centralized restore of an Always On availability group (AAG) database

Single or multiple databases can be restored from the same SQL host and instance. The databases can be restored either to the original SQL host or to an alternate SQL host with the following requirements:

- The alternate host must be a SQL virtual machine.
- The Microsoft application agent software must be installed and configured on the alternate host.

Note: When the Microsoft application agent is installed to perform a restore of a SQL virtual machine backup to an alternate SQL Server host, ensure that the SQL host was not previously registered to Data Manager as an Application Direct host.

- A system database cannot be restored to an alternate host or SQL instance.

The centralized restore of multiple SQL databases supports the following use cases:

- Performing disaster recovery of the original SQL instance.
- Performing a restore rehearsal by restoring a SQL instance database to an alternate host to validate the backups.

Centralized restore of SQL Virtual Machine backups can be performed from the **Recovery > Assets > SQL** window in the Data Manager UI.



Restore to original database - Restore databases to the original database on the original SQL Server where the SQL VM is protected by the Data Manager VM application-aware policy, the SQL VM is running, and SQL Server is operational. Backup copies are used for operational restore and for Disaster Recovery purposes.

For a single database restore, the restore can be performed from a most recent copy, a specific copy, or from a point-in-time (PIT) copy. For a single database, tail log backup can be performed for the active portion of the disk.

A restore of multiple databases can be performed only from the most recent backup copy.

Copy and Point-in-Time selection - Single database restore supports copy and Point-in-Time (PIT). With copy selection, FULL or TLOG can be selected and with PIT selection, roll-forward time can be selected and available only when TLOG backup is selected.



Note: Roll-forward time should be selected in between the timelines of full backup copy and log copy.

Restore to alternate database - Restore database to an alternate SQL Server where the SQL VM may (or may not) be protected by Data Manager. The Data Manager user can install agents for "unprotected" VM.

Data Manager provides a rich set of options for restoring SQL databases to alternate locations:

- Can select a target host
- Can select a target SQL Instance
- Can select a target database



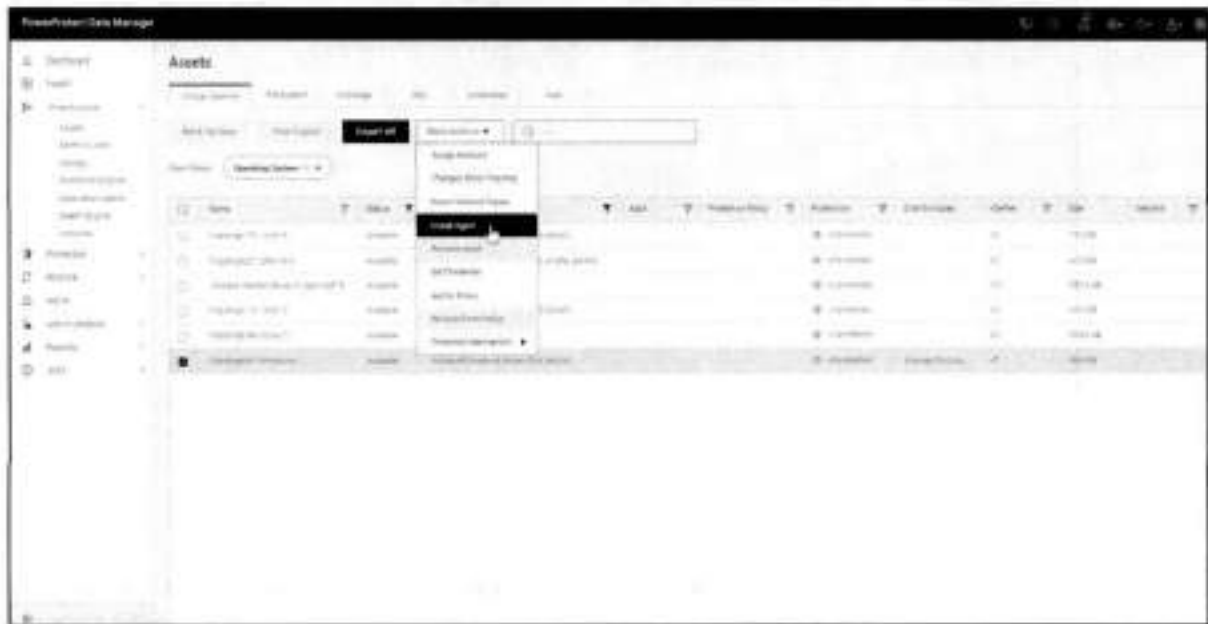
Refresh option to discover recently added SQL instances and databases on the host.

The target host must have an application agent installed:

- Pushed when VM is added to an application-aware protection policy
- Manually installed by a Data Manager administrator

The **Install Agent** option is available from the **More Actions** menu in **Infrastructure->Assets->Virtual Machines**. This is available only for Windows VM.

Note: Operating system credentials are required to install the Data Manager SQL Server Agent for Windows. Set the credential from **More Actions > Set Credential**.



Select **Install** to install the agent on the Windows VM. You can monitor the installation status from the Jobs window.



Data Manager will automatically discover the target hosts, SQL instances, and SQL databases.

Database File Location: You can select one of the following restore options to specify the file system location where the databases are restored:



- **Restore database files to the original file location (location at backup time)** - Restores the backup data to the file directory that was used during the backup and overwrites the existing contents.
- **Note:** If the directory path cannot be created during the centralized restore, the restore fails.
- **Restore database files to the default file location as set by SQL** - Restores the backup data to the default file directory as used by the SQL Server.
- **Restore database files to a user-specified file location** - Restores the backup data to file directories that are specified for the database files and log files.



Recovery state selection:

You can select one of the following options:

- **RESTORE WITH RECOVERY:** Leaves the database ready to use by rolling back the uncommitted transactions. Additional transaction logs cannot be restored.
- **RESTORE WITH NORECOVERY:** Leaves the database non-operational and does not roll back the uncommitted transactions. Additional transaction logs can be restored.

See [PowerProtect Data Manager Microsoft Application Agent SQL Server User Guide](#) for more details.

Self-service restore of SQL virtual machine backups

SQL databases that are backed up with an application aware VM protection policy can be restored using the Microsoft application agent.

Full and transaction log backups created by a PowerProtect virtual machine application-aware protection policy can be restored using Microsoft application agent tools. The backups are restored to a SQL Server hosted on a VMware virtual machine.

When a SQL Server virtual machine asset is added to a Data Manager virtual machine application aware protection policy, the Microsoft application agent and ItemPoint are silently installed on the protected SQL Server.

The Microsoft application agent automatically stores the DD series appliance host and login information from the protection settings that are configured in the Data Manager protection policy. This automatic configuration occurs when the SQL Server virtual machine asset is added to the Data Manager protection policy.

Restore operations can be performed using the Microsoft app agent for VM Direct SQL Server Management Studio (SSMS) plug-in or the command prompt. T-SQL scripts are not supported with VM Direct.

The Microsoft application agent can perform a database restore, table-level restore, or database instant access restore to the source virtual machine or to an alternate virtual machine. To perform restores to an alternate virtual machine, the destination virtual machine must be an asset of Data Manager.

However, instance-level restores can only be performed to the original source instance.

The Microsoft application agent supports both full backups and transaction log backups for a Microsoft Always On availability group (AAG). The AAG databases are indexed against the AAG cluster name. Full backups index the AAG database for all the AAG cluster nodes for one cycle of backup. Transaction log backups occur only on the preferred node.

Support for existing SQL agent backups with Data Manager

The Microsoft application agent enables onboarding existing stand-alone deployments, including their existing backups, to Data Manager. Existing backups are Microsoft application agent backups that are performed before integrating the Microsoft application

agent with the Data Manager software and added as an asset to a Data Manager protection policy.

With the onboarding capability, Data Manager provides the following centralized features:

- Visibility of both existing backups and any new self-service or Data Manager policy-driven backups.
- Automatic configuration of target protection storage based on the Data Manager protection policies that are used for database.
- All the other functionality that is provided for Data Manager protection policies.

Notes:

- After discovery is executed, it takes about 30 minutes for the copies to appear in the Data Manager under assets.
 - Onboarding of SQL backup copies to Data Manager is supported only from backups that are performed with Microsoft application agent 4.7 and later.
 - Up to three previous months of existing backups can be onboarded.
 - Retention lock is not supported for discovered existing backups in Data Manager.
 - Onboarding is not supported for DD Boost-over-FC backups and 32-bit FCI instance backups.
-

Disaster recovery

When a disaster scenario occurs, the Microsoft application agent supports disaster recovery of data on DD series appliances. The Microsoft application agent for Application Direct supports disaster recovery.

Procedure

1. Create a target Windows host with the same name as the source hostname.
2. Install a SQL Server instance with the same name as the source instance name.
3. Install the Microsoft application agent on the target Windows host.
4. Browse the backups of the source instance by selecting the appropriate storage unit.
5. Restore the system databases to the target instance.
6. Restore all the user databases to the target instance.

Technical support and resources

Dell.com/support is focused on meeting customer needs with proven services and support.

The [Data Protection Info Hub](#) provides expertise that helps to ensure customer success with Dell data protection products.

For more information, see the following related resources:

- **Product documentation:**
 - [PowerProtect Data Manager Microsoft Application Agent SQL Server User Guide](#)
 - [PowerProtect Data Manager Administration and User Guide](#)
 - **PowerProtect Data Manager E-LAB Navigator**—Provides compatibility information, including specific software and hardware configurations that PowerProtect Data Manager supports. To access E-LAB Navigator, go to [PowerProtect Data Manager Compatibility Matrix](#)

Dell EMC Integrated Data Protection Appliance with PowerProtect Data Manager Best Practices

Abstract

This document discusses best practices for using the Dell EMC™ Integrated Data Protection Appliance (IDPA) with Dell EMC PowerProtect Data Manager. This solution enables efficient and comprehensive data protection for proven and modernized workloads.

July 2020

Revisions

Revisions

Date	Description
July 2020	Initial release

Acknowledgments

Author: Sandeep Rajagopal

The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [7/29/2020] [Best Practices] [H18448]

Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents.....	3
Executive summary.....	4
1 Introduction.....	5
1.1 Integrated Data Protection Appliance.....	5
1.1 PowerProtect Data Manager.....	5
2 PowerProtect Data Manager with IDPA overview.....	7
2.1 Example use case: Protecting Kubernetes workloads.....	7
2.2 Key benefits.....	8
3 Technical considerations and deployment.....	9
3.1 Prerequisites.....	9
3.2 PowerProtect Data Manager deployment.....	9
3.3 Data Protection Central deployment.....	10
3.3.1 Disabling Integrated Data Protection Appliance System Manager.....	10
3.3.2 Registering IDPA and PowerProtect Data Manager to an external instance of Data Protection Central.....	12
3.4 Discovering the embedded Data Domain system with PowerProtect Data Manager.....	15
4 Interoperability and limits.....	18
5 Summary.....	19
A Technical support and resources.....	20
A.1 Related resources.....	20

Executive summary

Data protection is an integral and essential part of any successful business. Organizations and IT teams require a proven, powerful, modern, scalable, and easy-to-use data protection solution.

To meet market demands, Dell Technologies™ offers a unique solution consisting of Dell EMC™ PowerProtect Data Manager with Dell EMC Integrated Data Protection Appliance (IDPA). The IDPA protects leading enterprise applications and operating systems. The PowerProtect Data Manager solution protects traditional workloads such as file system, SQL, SAP® HANA®, Oracle®, and Microsoft® Exchange. It also protects modern, cloud-native Kubernetes workloads and enables differentiated VMware® protection.

Integrating the proven IDPA with the modern PowerProtect Data Manager enables you to get the most efficient and comprehensive data protection solution for your modernized workloads. Also, you can use the Dell EMC Data Domain™ system that is bundled with IDPA as back-end storage for PowerProtect Data Manager workloads at no extra cost.

Audience

The information in this document is intended for customers who are responsible for planning, implementing, and administering the environments that contain IDPA solutions. The primary audience consists of customers, customer service, and remote Professional Services engineers.

1 Introduction

This section provides an overview of the Integrated Data Protection Appliance and PowerProtect Data Manager software.

1.1 Integrated Data Protection Appliance

The Integrated Data Protection Appliance (IDPA) is an all-in-one backup appliance that reduces the complexity of managing multiple data silos, point solutions, and vendor relationships. IDPA simplifies deployment and management while delivering powerful, enterprise-grade data protection capabilities for small, midsized, and enterprise organizations with a low cost-to-protect ratio.

The IDPA provides a solution for data-protection administrators who are accustomed to configuring and managing one or more data-protection and storage devices, but are challenged to manage independent and disconnected applications.

IDPA System Manager enables administrators to efficiently manage the IDPA components from a single user interface. This interface includes monitoring, reporting, analytics, and search capabilities to help simplify the data-protection experience.

The IDPA (Figure 1) streamlines the configuration and the integration of data-protection components in a consolidated solution and also offers the following benefits:

- Simplified deployment and configuration
- Backup administration
- Deduplication
- Native cloud data reduction (DR) and long-term retention (LTR)
- Instant access and restore
- Monitoring and analytics
- Search
- Scalability
- Unified support



Figure 1 Dell EMC Integrated Data Protection Appliance

1.1 PowerProtect Data Manager

Dell EMC 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 enables the transformation from traditional centralized protection to an IT-as-a-service model that is based on a self-service design. This design ensures that you can enforce compliance and other business rules, even when backup responsibilities are decentralized to individual database administrators and application administrators.

PowerProtect Data Manager (shown in Figure 2) supports multiple workloads such as the following:

- Kubernetes
- Filesystem
- Databases
- Virtual machines
- Storage-Direct (SDM)



Figure 2 PowerProtect Data Manager

2 PowerProtect Data Manager with IDPA overview

The IDPA protects leading enterprise applications and operating systems. The PowerProtect Data Manager solution protects traditional workloads such as file system, SQL, SAP® HANA®, Oracle®, and Microsoft® Exchange. It also protects modern, cloud-native Kubernetes workloads and offers differentiated VMware® protection.

Integrating the proven IDPA with modern PowerProtect Data Manager enables you to get the most efficient and comprehensive data protection solution for your modernized workloads. Also, you can use the Dell EMC Data Domain™ system that is bundled with IDPA as back-end storage for PowerProtect Data Manager workloads at no extra cost.

PowerProtect Data Manager delivers a simplified user interface that is easy to navigate. It eases the process and reduces the number steps required for creating backups, replication, recovery, expansion, and upgrades. Monitoring, managing, and analyzing copies of data are no longer tedious tasks with PowerProtect Data Manager, which also addresses the issue of copy sprawling.

PowerProtect Data Manager provides centralized governance that helps mitigate risk and assures compliance of service-level agreements (SLAs) and service-level objectives (SLOs) through simple protection workflows.

PowerProtect Data Manager enables automated discovery and onboarding of the following:

- Databases
- Virtual machines
- Kubernetes clusters
- Data Domain protection storage

It also offers self-service and centralized protection for Microsoft® SQL Server®, Microsoft Exchange, SAP HANA, and Oracle databases, and is multicloud optimized for efficient, long-term retention and disaster recovery.

When you use the embedded Data Domain in IDPA as the target storage for unique workloads, you get secondary storage with unmatched efficiency, deduplication, performance, and scalability at no extra cost. This result combines the best of proven and modern data protection solutions.

2.1 Example use case: Protecting Kubernetes workloads

Today, containers have increased in popularity. Containers are like virtual machines but have relaxed isolation properties to share the operating system. This container has its own file system, CPU, memory, and process space. Key benefits of containers include agile application creation, continuous development, environmental consistency across development, application-centric management, efficient resource allocation, and resource isolation.

With distributed container deployment, it is necessary to protect the workloads. PowerProtect Data Manager allows protecting the production workloads in Kubernetes (K8s) environments. This capability ensures the data is easy to back up and restore, always available, consistent, and durable in a Kubernetes workload or DR situation.

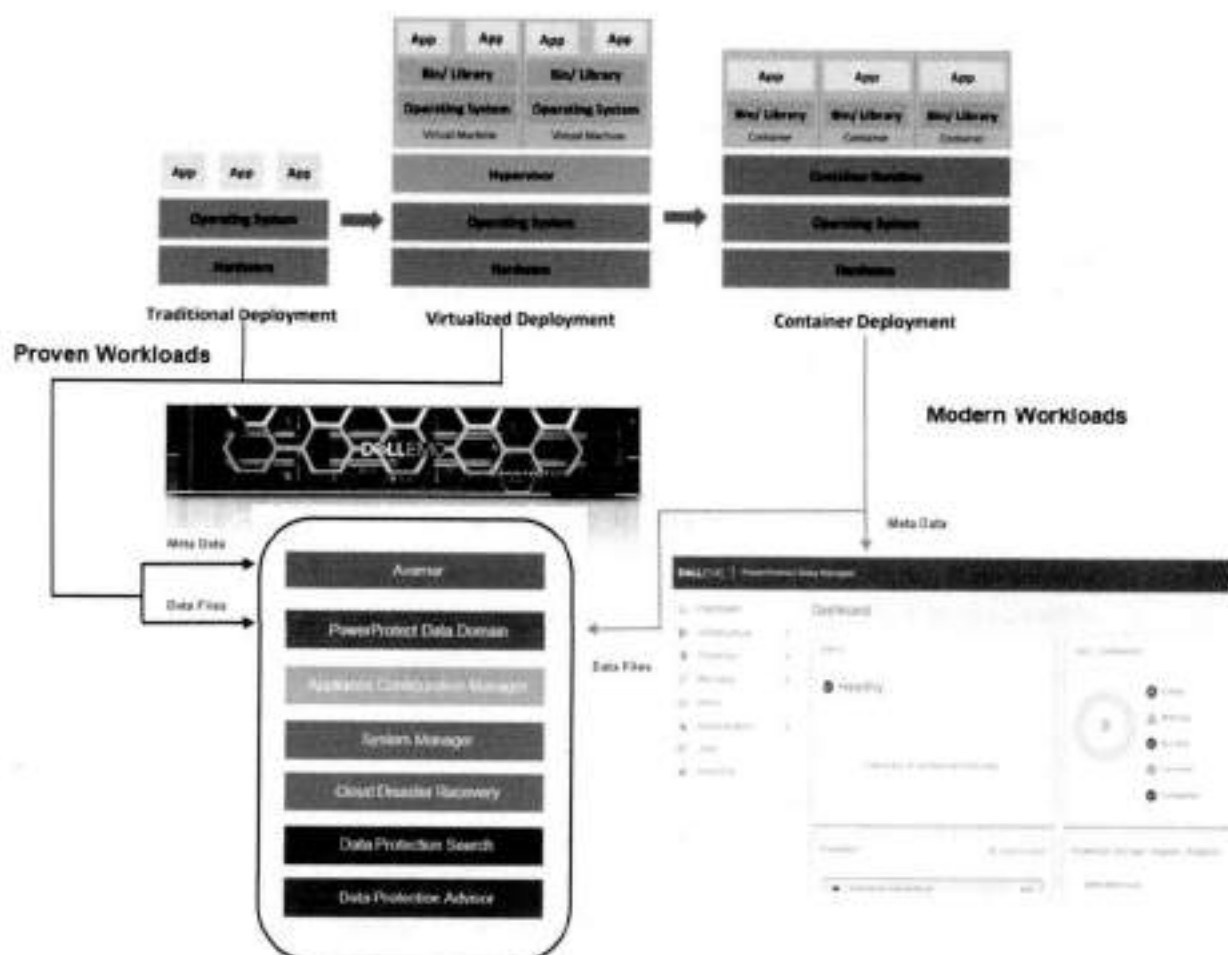


Figure 3 PowerProtect Data Manager with IDPA

2.2 Key benefits

With IDPA protecting leading enterprise applications, virtual machines, and operating systems, and with PowerProtect Data Manager protecting modern workloads such as Kubernetes, you have an efficient and optimized data-protection solution.

For modern workloads, you are protected with embedded Data Domain systems, and can have secondary storage with unmatched efficiency, deduplication, performance, and scalability at no extra cost.

For more information about configuring PowerProtect Data Manager, see the [PowerProtect Data Manager Administration and User Guide](#).

3 Technical considerations and deployment

This section provides prerequisites, best practices, and deployment guidance for using IDPA with PowerProtect Data Manager.

3.1 Prerequisites

Ensure that Dell EMC IDPA is deployed, configured, and operating.

3.2 PowerProtect Data Manager deployment

It is vital to plan for the environment that is used for deploying PowerProtect Data Manager, and to facilitate adequate resources to achieve optimal performance of PowerProtect Data Manager.

The PowerProtect Data Manager software appliance is easy to install and configure. You can deploy the PowerProtect Data Manager Open Virtual Appliance (OVA) using one of the following methods:

- Manually deploying the OVA to a VMware vCenter® server: Use this method to deploy the OVA to a stand-alone or cluster host while logged into the vCenter server. This method allows you to configure the network settings during deployment.
- Manually deploying the OVA to a VMware ESXi™ host: Use this method to deploy the OVA while logged in to an ESXi host. Use the VM console to configure the network settings after the deployment completes.

Minimum requirements:

The minimum resource requirements to deploy a PowerProtect Data Manager in vSphere 6.0 and above are as follows:

- 10 CPU cores
- 18 GB RAM
- 700 GB disk space

If you plan to use **Cloud DR**, your system must also meet the following requirements:

- 14 CPU cores
- 22 GB RAM

See section 4 to check for interoperability.

Best practices:

- Deploy the PowerProtect Data Manager Software on an external vCenter Server.
- Create a dedicated PowerProtect vCenter user and avoid using vCenter administrator credentials.
- Create a dedicated PowerProtect Data Domain BOOST User for Data Domain Discovery.

See the *PowerProtect Data Manager Deployment Guide* for more deployment information.

3.3 Data Protection Central deployment

It is critical to plan for the environment that is used to deploy Data Protection Central and provide the appropriate resources to optimize its performance. This section describes how to install and configure Data Protection Central.

Before you begin deployment, review the following information:

- Ensure the network is set up with IDPA and PowerProtect Data Manager is set up.
- Ensure the DNS is set up correctly. The correct DNS setup ensures that systems monitored by Data Protection Central can resolve the Data Protection Central hostname and Fully Qualified Domain Name (FQDN).
- Ensure the time synchronization between the IDPA and PowerProtect Data Manager is correct for efficient and effective monitoring of components.
- Deploy the Data Protection Central Open Virtualization Appliance (OVA) using a VMware vSphere® client.
- Ensure the minimum resource requirements are met to deploy Data Protection Central:
 - 4 CPU cores
 - 8 GB RAM
 - 550 GB disk space

Note: The Data Protection Central OVA does not deploy directly on an ESXi server.

- For optimal performance, you must deploy the external Data Protection Central component outside the IDPA vCenter for optimal performance.
- See section 4 to check for interoperability.

For more information, see the *Data Protection Central Getting Started Guide*.

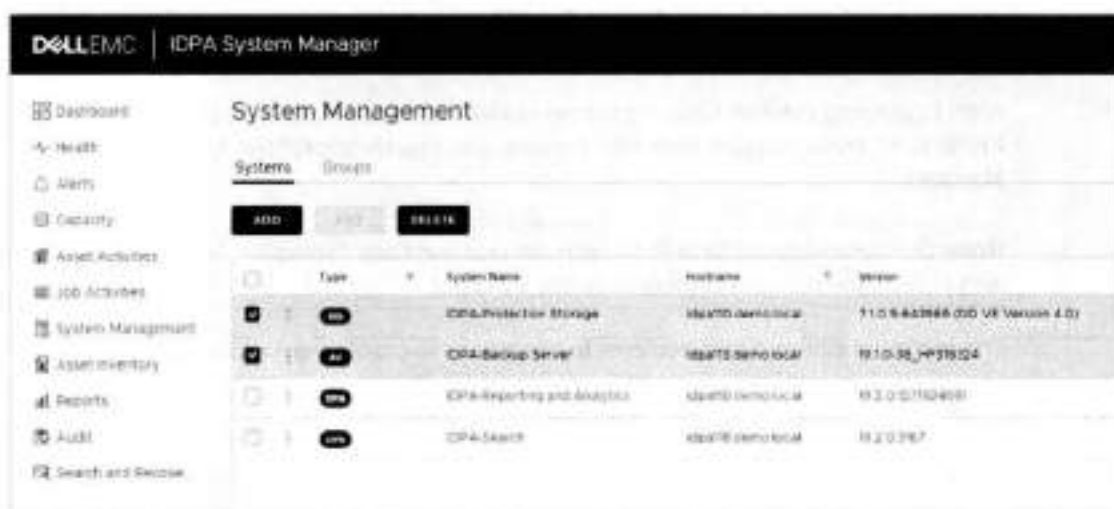
3.3.1 Disabling Integrated Data Protection Appliance System Manager

IDPA System Manager enables administrators to efficiently monitor and manage the software products within the IDPA from a single user interface, simplifying the data-protection experience.

With PowerProtect Data Manager and IDPA, you cannot use the embedded system manager to efficiently monitor the PowerProtect Data Manager activities. To efficiently monitor both IDPA and PowerProtect Data Manager, you must deploy an external instance of Data Protection Central.

To unregister the IDPA backup server and IDPA protection storage components from the embedded system manager, perform the following steps.

1. Log in to **IDPA System Manager**.
2. In the left menu, click **System Management**.
3. Select **IDPA-Backup Server** and **IDPA-Protection Storage**.



4. Click **Delete**. The **Confirm Delete** window appears.



5. Click **Delete**.

The system is removed. A deactivation activity message appears on the **Audit** page.

Note: Follow the above procedure to unregister Data Protection Advisor and Data Protection Search from the embedded System Manager.

3.3.2 Registering IDPA and PowerProtect Data Manager to an external instance of Data Protection Central

With PowerProtect Data Manager and IDPA, it can be cumbersome to manage and monitor the activities of a proven and modern data-protection solution individually.

Deploying an external Data Protection Central component enables administrators to efficiently monitor and manage both IDPA and PowerProtect Data Manager from a single user interface. This ability simplifies the entire data protection experience for the customer.

Also, registering external Data Protection Advisor and Data Protection Search components with Data Protection Central enables common reporting and search capabilities for both IDPA and PowerProtect Data Manager.

Note: Only one external Data Protection Advisor and Data Protection Search component can be configured with an external instance of Data Protection Central.

The minimum software requirements to complete this process are as follows:

- PowerProtect Data Manager 19.5
- Data Protection Central 19.3
- Data Protection Advisor 19.3 (external)
- Data Protection Search 18.2/19.1/19.2/19.3 (depending on IDPA models)

This process assumes that an external Data Protection Central instance has been deployed.

Adding an IDPA backup server:

1. Log in into Data Protection Central.
2. In the left menu, click **System Management**.
3. Click **Add**. The **Add System** window is displayed.
4. On the **Select System Type** page, select **Avamar**, and click **Next**.

5. On the **Connection Information** page, specify the following information:
 - **Name:** Specify a name to identify the system.
 - **Hostname:** Specify the FQDN of the Avamar system.
 - **Avamar Username:** Specify the username of the Avamar system. For the Avamar administrator, the username is **MCUser**.
 - **Avamar Password:** Specify the password for the Avamar system user interface.

Note: Operating-system root credentials are optional for Avamar 19.3 and later. If the Avamar version is **19.2 or earlier**, click the **toggle** button to enable the field and specify the **OS root > password**.

The screenshot shows a window titled "Add System" with a sidebar on the left containing three steps: "1 Select System Type", "2 Connection Information" (which is selected and highlighted), and "3 Certificate Verification". The main area is titled "Connection Information" and contains the following fields and controls:

- Name:** IDPA-Backup Server
- Hostname:** idpa10.demo.local
- Avamar Username:** MCUser
- Avamar Password:** A masked password field with a toggle icon to the right.
- Version 19.2 or lower:** A toggle switch that is currently turned off.
- OS Root Password:** A masked password field with a toggle icon to the right.
- Show optional fields:** A toggle switch that is currently turned off.
- Buttons:** CANCEL, BACK, and NEXT.

6. Click **Next**.
7. Select the **Accept Certificate** option and click **Save**.

Note: Allow 15 minutes for the data to be synchronized between the Data Protection Central and IDPA backup server.

Adding IDPA protection storage:

1. Log in into Data Protection Central.
2. In the left menu, click **System Management**.
3. Click **Add**. The **Add System** window is displayed.
4. On the **Select System Type** page, select **Data Domain**, and click **Next**.

5. On the **Connection Information** page, specify the following information:
 - **Name:** Specify a name that identifies the system.
 - **Hostname:** Specify the FQDN of the Data Domain system.
 - **Username:** Specify the Data Domain administrator username.
 - **Password:** Specify the Data Domain administrator password.

The screenshot shows a window titled "Add System" with a close button (X) in the top right corner. On the left side, there is a vertical progress bar with three steps: "1 Select System Type", "2 Connection Information" (which is highlighted), and "3 Certificate verification". The main area of the window is titled "Connection Information" and contains the following fields:

Type	Data Domain
Name	IDPA-Protection Storage
Hostname	ipafid.demo.local
Username	sysadmin
Password	*****

At the bottom right of the window, there are three buttons: "CANCEL", "BACK", and "NEXT".

6. Click **Next**.
7. Select the **Accept Certificate** option and click **Save**.

Note: Allow 15 minutes for the data to be synchronized between the Data Protection Central and IDPA protection storage.

Adding PowerProtect Data Manager:

1. Log in into Data Protection Central.
2. In the left menu, click **System Management**.
3. Click **Add**. The **Add System** window is displayed.
4. On the **Select System Type** page, select **PowerProtect Data Manager**, and click **Next**.

5. On the **Connection Information** page, specify the following information:
 - **Name:** Specify a name that identifies the system.
 - **Hostname:** Specify the FQDN of the PowerProtect system.
 - **Username:** Specify the PowerProtect administrator username.
 - **Password:** Specify the PowerProtect administrator password.

Field	Value
Type	PowerProtect Data Manager
Name	PowerProtect Data Man
Hostname	ppdm-demo.local
Username	administrator
Password	*****

6. Click **Next**.
7. Select the **Accept Certificate** option and click **Save**.

Note: You can also integrate Data Protection Advisor and Data Protection Search to the external Data Protection Central instance using methods similar to the above process.

See the *Data Protection Central Administrator Guide* for more information.

3.4 Discovering the embedded Data Domain system with PowerProtect Data Manager

PowerProtect Data Manager leverages the embedded Data Domain in the IDPA as the target server for protecting modern workloads. To start protecting modern workloads, PowerProtect Data Manager must discover the embedded Data Domain system as the target storage server. As a prerequisite, create a dedicated PowerProtect Data Domain Boost User for Data Domain Discovery.

Discovering the embedded Data Domain System:

1. Log in to PowerProtect Data Manager.
2. Click **Infrastructure > Storage**. The **Storage** window appears.
3. In the **Protection Storage** tab, click **Add**.
4. In the **Add Storage** dialog box, select the **Data Domain** system.

5. Specify the **Embedded Data Domain** system attributes:
 - **Name:** Specify the embedded Data Domain name.
 - **Address:** Specify the hostname, FQDN, or the IP address.
 - **Port:** Specify the port for SSL communication. The default value is 3009.

The screenshot shows the 'Add Storage' dialog box with the following fields and values:

- Name:** DPA-Prebuilt Storage
- Address:** host10.smb.local
- Port:** 3009
- Host Credentials:** DATA DOMAIN - (dropdown menu open)
- Certificate:** VERIFY (button highlighted in black)

6. Under **Host Credentials**, click **Add**.
7. Enter the following information in the **Add Credentials** window.
 - **Name:** Provide a descriptive name to the keychain.
 - **Username:** Specify the Data Domain Boost Username for the Data Domain Discovery.
 - **Password:** Specify the Data Domain Boost Credentials.

The screenshot shows the 'Add Credentials' dialog box with the following fields and values:

- Type:** DATA DOMAIN (dropdown menu)
- Name:** DPA-Data Domain
- User Name:** ddboost_admin
- Password:** [Masked]

8. Click **Save**.
9. Click **Verify** to review the certificate, and click **Accept**.

Technical considerations and deployment

10. Click **Save** to exit the **Add Storage** window and initiate the discovery for the **embedded Data Domain System**.

A dialog box appears to indicate that the request to add storage has been initiated.

11. In the **Storage** window, click **Discover** to refresh the window with the embedded Data Domain system information.
12. When a discovery completes successfully, the Status column updates to **OK**.

This process has successfully discovered the embedded Data Domain system as the target storage server with PowerProtect Data Manager. You can now start protecting modern workloads in the embedded Data Domain system, benefiting from secondary storage with unmatched efficiency, deduplication, performance, and scalability.

4 Interoperability and limits

The PowerProtect Data Manager with IDPA solution has the following the minimum software requirements:

- PowerProtect Data Manager 19.5
- Data Protection Central 19.3 (external)
- Data Protection Advisor 19.3 (external)
- Data Protection Search 18.2/19.1/19.2/19.3 (depending on IDPA models)

According to the architecture of PowerProtect Data Manager, for each policy, there is a Storage-Unit that is created on the embedded Data Domain system. You must ensure that you do not run out of Storage-Unit space on the embedded Data Domain system.

Figure 4 (from [E-Lab Navigator](#)) shows the maximum number of MTrees that can be created on the embedded Data Domain system and the total number of MTrees that can be active for both read and write operation.

	Internal IDPA Software					External to IDPA		
	DD/VE version	DD OS version	Mtree limits Total/Active	AV version	DPS version (can be external)	PPDM version	External DPC version	External IDPA version
	IDPA 2.3x							
DP4x	4.0	6.2.0	100/32	18.2	18.2/19.1	19.5	19.3	19.3
DP5x	NA	6.2.0	128/128	18.2	18.2/19.1	19.5	19.3	19.3
DP6x	NA	6.2.0	256/256	18.2	18.2/19.1	19.5	19.3	19.3
	IDPA 2.4x							
DP4x	4.0	6.2.0	100/32	18.2	18.2/19.1	19.5	19.3	19.3
DP5x	NA	6.2.0	128/128	18.2	18.2/19.1	19.5	19.3	19.3
DP6x	NA	6.2.0	256/256	18.2	18.2/19.1	19.5	19.3	19.3
	IDPA 2.5							
DP4x	4.0	7.1.0.5	100/32	19.1	19.2	19.5	19.3	19.3
DP5x	NA	7.1.0.5	128/128	19.1	19.2	19.5	19.3	19.3
DP6x	NA	7.1.0.5	256/256	19.1	19.2	19.5	19.3	19.3
	IDPA 2.6							
DP4x	5.0	7.2.x	100/32	19.3	19.3	19.5	19.3	19.3
DP5x	NA	7.2.x	128/128	19.3	19.3	19.5	19.3	19.3
DP6x	NA	7.2.x	256/256	19.3	19.3	19.5	19.3	19.3

Figure 4 Interoperability and limits

Note: One MTree is always reserved for the IDPA backup server.

5 Summary

Dell EMC PowerProtect Data Manager with the Dell EMC IDPA backup appliance provides the benefits of both proven and modern data-protection solutions for modernized workloads. With protection for unique workloads in the embedded Data Domain systems, you can also benefit from secondary storage with unmatched efficiency, deduplication, performance, and scalability at no extra cost.

Also, the external Data Protection Central component enables administrators to efficiently monitor and manage both the IDPA and PowerProtect Data Manager from a single user interface, simplifying the entire data protection experience.

A Technical support and resources

Dell.com/support is focused on meeting customer needs with proven services and support.

[Storage technical white papers and videos](#) provide expertise that helps to ensure customer success on Dell EMC storage platforms.

A.1 Related resources

For additional information, see the following resources:

- [PowerProtect Data Manager Administrator and User Guide](#)
- [PowerProtect Data Manager Deployment Guide](#)
- [Data Protection Central Getting Started Guide](#)
- [Data Protection Central Administrator Guide](#)
- [E-Lab Navigator](#): Provides compatibility information, including specific software and hardware configurations that the solution supports.

Dell EMC PowerProtect Data Manager: Microsoft Exchange Backup and Recovery

December 2021

H18560.1

White Paper

Abstract

This white paper focuses on protecting a Microsoft Exchange database using Dell EMC PowerProtect Data Manager, the next-generation data protection platform.

Dell Technologies

Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2020-2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA December 2021 H18560.1.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Contents

Executive summary.....	4
Overview.....	5
Supported configurations.....	8
Installation and configuration of Microsoft application agent for Exchange Server.....	9
Microsoft application agent for Exchange database backup.....	13
Restoring Exchange Server databases.....	20
Cloud tiering and replication.....	24
Disaster recovery of Exchange Server.....	24
Conclusion.....	26
Technical support and resources.....	26

Executive summary

Introduction

Today's data protection is either too complex, requires multiple vendors, does not scale, or fails to meet the needs of fast-growing, modern, and agile organizations. As businesses continue to consume IT resources differently, there is a need for powerful, efficient, and trusted data protection. These solutions can enable organizations to transform and meet future demands when modernizing their IT environments.

Organizations strive to provide users with large mailboxes while reducing the requirements and complexity of their business-exchange backup data storage. With current trends, the user mailbox size is growing rapidly. This growth makes it more challenging, if not impossible, to back up all the business-exchange data within the nightly backup window.

Dell EMC PowerProtect software is the next-generation data management platform that transforms traditional data protection into comprehensive data management. PowerProtect software is defined with integrated deduplication for data protection, replication, and reuse.

This white paper outlines Microsoft Exchange protection and recovery with PowerProtect software, which provides reliable and efficient data protection functionalities. The PowerProtect Microsoft application agent uses block-based backup technology to back up Exchange Server databases in stand-alone and database availability group (DAG) environments. This block-based technology tracks the changed blocks of the Exchange database and log files. A full backup backs up each selected Exchange database and its log files. An incremental backup backs up only the changed blocks.

Block-based backups are fast backups that have reduced backup times. This advantage is due to the way that the backup process respectively backs up only the occupied disk blocks and changed disk blocks of the Exchange database and log files. Block-based backups provide instant access to the backups.

Audience

This white paper is intended for customers, partners, and employees who want to better understand, evaluate, and explore Dell EMC PowerProtect software for Exchange server backup and recovery.

Revisions

Table 1. Revisions

Date	Description
October 2020	Initial release
December 2021	Template update

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#) (subject line: Feedback for document: H18560.1).

Author: Sonali Dwivedi

Note: For links to other documentation for this topic, see the [Data Protection Info Hub](#).

Overview

Introduction

The Microsoft application agent enables an application administrator to protect and recover Exchange application data on the application host. Data Manager integrates with the Microsoft application agent to check and monitor backup compliance against protection policies. Data Manager also enables centralized scheduling for backups.

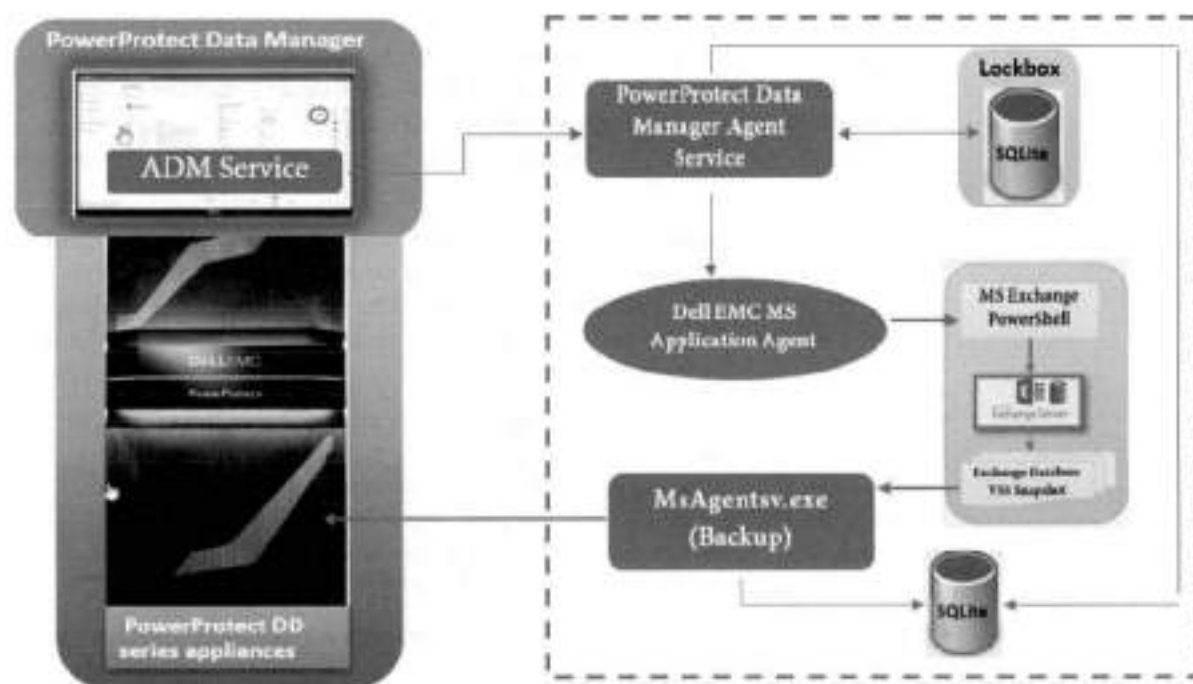


Figure 1. Microsoft application agent for Exchange Server

Protection storage

The first step in configuring Exchange Server integration with Data Manager is to ensure that you have a protection storage target that is configured to store the backup data. You can accomplish this task by adding one or multiple PowerProtect DD series appliances and add the protection storage using the Data Manager UI.

PowerProtect Data Manager

Use 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 are produced by the application agent, and monitor that catalog data to determine if retention policies are being adhered to
- Manage the life cycle of backups that are created by the application agent (ensure that the backups are marked for garbage collection based on the rules of the retention policy)

Data Manager does not change the way that the application agent works. Database administrators or backup administrators create the backups and perform the restores.

PowerProtect Data Manager agent service

The PowerProtect agent service is a REST API-based service, and the application agent installs this service on the application host. The agent service provides services and APIs for discovery, protection, restore, instant access, and other related operations. The Data Manager uses the agent service to provide integrated data protection for the application assets. The PowerProtect agent service provides important functionality for the application agent operations with the Data Manager. The PowerProtect agent service performs the following operations:

- **Add-on detection:** An add-on integrates the application agent into the agent service. The agent service automatically detects the add-ons on the system for each application asset type and notifies the PowerProtect Data Manager. While multiple add-ons can operate with different asset types, only one agent service runs on the application host. Specific asset types can co-exist 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 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 Data Manager. When you add an asset to a protection policy for self-service or centralized protection, the Data Manager automatically pushes the protection configuration to the agents. The Data Manager also performs this action if you modify the protection policy, including changing the DD Boost credentials.
- **Centralized backups:** The agent service performs the centralized backups as requested by the Data Manager.
- **Centralized restores:** The agent service performs the centralized restores as requested by the Data Manager.

Exchange PowerShell interface for the application agent

The Exchange Management Shell is built on Microsoft Windows PowerShell technology and provides a powerful command-line interface that enables the automation of Exchange administration tasks. You can use the Exchange Management Shell to perform every task in the Exchange graphical management tools, plus tasks that you can perform elsewhere (for example, bulk operations). The application agent has PS/cmdlet interface that blends in with the Exchange shell, which gives Exchange administrators a seamless experience.

SQLite database

SQLite is a C library that provides a lightweight disk-based database that does not require a separate server process. This library allows you to access the database using a nonstandard variant of the SQL query language. Some applications can use SQLite for internal data storage. This library is used in this solution to store information about all types of backups like self-service, centralized, and automatic log backup for Data Manager.

Application Discovery Manager

The PowerProtect Data Manager Application Discovery Manager (ADM) provides continuous discovery and mapping of applications. It also maps their dependencies and configurations concerning their underlying infrastructure in data-center environments. ADM allows accurate, real-time visibility into the data center from an application standpoint. It is critical for planning data-center consolidations and migrations as well as managing change impact, virtualization initiatives, and disaster recovery.

Lockbox

The lockbox is an encrypted file that the Microsoft application agent uses to store confidential data, such as login credentials, and protect that data from unauthorized access. For each PowerProtect protection policy, the Data Manager creates a storage unit and automatically configures the lockbox on the application host. A source lockbox and replication target lockbox are created and configured on the application host.

When you first use the Data Manager UI to add the Microsoft application agent and create the protection policy for Exchange data protection, Data Manager automatically configures the lockbox for the Exchange server. The lockbox for the Microsoft application agent is created in the default directory `C:\ProgramFiles\DPSAPPS\common\lockbox`. Data Manager integration requires the lockbox to be in the default directory.

Support for existing Microsoft application agent backups

The Microsoft application agent provides the capability to onboard existing stand-alone deployments, including their existing backups, to Data Manager. Existing backups are Microsoft application agent backups that are performed before you integrate the Microsoft application agent with the Data Manager software. They are also performed before you add an asset to a Data Manager protection policy. Note the following information regarding existing backups:

- Onboarding of Exchange backup copies to Data Manager is supported only from backups performed with Microsoft application agent version 4.7 and later.
- You can onboard up to three previous months' worth of existing backups.
- Retention lock is not supported for discovered existing backups in Data Manager.
- Onboarding of DD Boost-over-FC backups is not supported.

With the onboarding capability, Data Manager provides the following centralized features:

- Visibility of both existing backups and any new self-service or Data Manager policy-driven backups of onboarded assets.
- Automatic configuration of target protection storage based on the Data Manager protection policies that are used for your database.
- All other functionality that is provided for Data Manager protection policies.
- Ability to create a storage unit (when creating a protection policy) on the specified DD system backup host that is managed by Data Manager. All subsequent backups of assets in that protection policy go to this new storage unit. This implementation takes the storage-unit information that is provided before you onboard triggering backups through scripts and overrides it with the storage-unit information that is provided by Data Manager.

Supported configurations

The solution described in this document supports the following configurations.

- PowerProtect Data Manager supports the coexistence of the Microsoft application agent and the File System agent on Windows.
- The Microsoft application agent does not support the Meta-Cache Database (MCDB) feature in Exchange Server 2019. Do not enable MCDB in Exchange Server 2019.

Table 2. Supported configurations

Category	Feature	PowerProtect support for Microsoft application agent
Configurations	IP DAG	Yes
	IP-LESS DAG	Yes
	Parent Child Domain	No
	Disjoint namespace	No
	Standalone	Yes
	Dual NIC	Yes
Backup	Full and Incremental	Yes
	Retention Management	Yes
	IP-Less DAG backup	Yes
	IP DAG	Yes
	Log truncation	Yes
	Application Consistent	Yes
	Writer level backup and database level	Yes
	Roll over Parallelism	No
	Exclude components during backup	Yes
Restore	Point in Time Restore	Yes
	Flat file Restore	No
	Recover to alternate Database	Yes
	GLR To and from RDB	No
	GLR to Alternate mailbox	Yes
	GLR to PST files	Yes
	GLR to non-Exchange Server	No
	GLR to Messages and Text	Yes

Category	Feature	PowerProtect support for Microsoft application agent
	Proxy GLR	No
	Roll Forward Recovery	No
	Redirected Restore	Yes
Backup and restore technology		VSS for Snapshot Block Based Backup for CBT/Data Transfer
PowerProtect DD series appliance support	PowerProtect DD series appliance	Yes
	PowerProtect DD Virtual Edition	Yes
	PowerProtect DD Retention lock	Yes
	PowerProtect DD cloud tier for LTR	Yes
Cloud support	Microsoft Azure	No
	Amazon Web Services (AWS)	No
	ECS	Yes
Exchange versions supported	Exchange Server 2010	Yes
	Exchange Server 2013	Yes
	Exchange Server 2016	Yes
	Exchange Server 2019	Yes
	Windows Server 2012	Yes
	Windows Server 2012 R2	Yes
	Windows Server 2016	Yes
	Windows Server 2019	Yes
	Windows Server Core 2019	No
Cloning		Yes

Installation and configuration of Microsoft application agent for Exchange Server

Introduction

The Microsoft application agent enables an application administrator to protect and recover the Exchange application data on the application host. Data Manager integrates with the Microsoft application agent to check and monitor backup compliance against protection policies. Data Manager also enables central scheduling for backups.

Prerequisites

Ensure that your environment meets the following requirements for a new deployment or upgrade of Data Manager:

- A list of hosts that write backups to DD systems is available.
- DD OS version 6.1 or later and the DD Management Console (DDMC). All models of DD series systems are supported.

Note: DDMC is required with a DD OS version earlier than 6.1.2. With DD OS version 6.1.2 or later, you can add and use a DD series system directly without DDMC.

- Application agent version 19.5 or later is installed.
- License: A trial license is provided with the PowerProtect Data Manager software. For DPS applications, backup, and enterprise, you can contact Dell EMC Licensing Support for assistance with a permanent Data Manager license.
- Large environments require multiple Data Manager instances. Contact Champions eCDM@emc.com for assistance with sizing requests.
- The PowerProtect Data Manager 19.6 download file requires:
 - VMware ESXi version 6.0, 6.5, 6.7, or 7.0
 - 8 vCPUs, 18 GB RAM, one 100 GB disk, and one 500 GB disk
 - The latest version of the Google Chrome browser to access the Data Manager UI
 - TCP port 7000 is open between Data Manager and the application agent hosts
- The Exchange Server environment meets the following prerequisites before you install the Microsoft application agent. Install the following applications on the Windows host:
 - Microsoft Exchange Server
 - .NET Framework 4.0 or later
- If installing ItemPoint for granular-level recovery, install .NET Framework 4.5.

Installation and configuration

Perform the following steps to install and configure the Microsoft application agent:

Prerequisite: Install PowerProtect DD Management Center (DDMC). Data Manager uses DDMC to connect to the DD systems. See the [PowerProtect DD Management Center Installation and Administration Guide](#) for instructions.

1. Install the Microsoft application agent on the Exchange Server host.
 - a. In the **Data Manager** UI, click **System Settings > Agent Downloads**, select the Microsoft application agent download package **msappagent196_win_x64.zip**, and download the package to the Windows Exchange Server host.
 - b. Log in to the Exchange Server host as an Administrator to install the Microsoft application agent.
2. Configure the required user privileges on the Exchange Server host using **App Agent Exchange Admin Configuration** tool.

3. Add or approve the Microsoft application agent in Data Manager.
4. Discover the Exchange application host.
5. Create a protection policy to protect the Exchange host.
 - On each node in the DAG, repeat the steps to install the Microsoft application agent, and add and discover the application host in Data Manager.
 - Protection of the nodes in a DAG requires that all the nodes be registered to the Data Manager server.
 - You cannot perform a backup to a secondary PowerProtect DD series appliance. You can only restore from a secondary PowerProtect DD series appliance.



User configuration: App Agent Exchange Admin Configuration tool

Note the following information when using the App Agent Exchange Admin Configuration tool:

- To protect a stand-alone Exchange Server or Exchange DAG with the Microsoft application agent, you must configure an account with the required privileges. The App Agent Exchange Admin Configuration tool simplifies configuring security group memberships by ensuring that users have all the required Active Directory security group memberships and PowerShell management roles.
- To use the App Agent Exchange Admin Configuration tool, you must be logged in with domain administrator permissions. You can use an existing non-administrative user to run the App Agent Exchange Admin Configuration tool. However, this action is only possible if you select **Skip Active Directory Authentication** and configure the user on each Exchange Server node. This option skips the Active Directory authentication and authorization operations for the user. It only sets the user as the Microsoft application agent Exchange user account in the registry for backup and recovery operations.
- The Microsoft application agent uses the user account that is set in the registry by the App Agent Exchange Admin Configuration tool to perform backups and database- or granular-level recovery.

- To create a Microsoft application agent Exchange administrator account, the App Agent Exchange Admin Configuration tool performs the following steps:
 - a. Creates an Active Directory user account
 - b. Creates the custom Exchange security group **Dell EMC App Agent Exchange Admin Roles**
 - c. (Optional) Allows selecting **Assign Organization Management** rights

Members of the Organization Management role group have permissions to manage Exchange objects and their properties in the Exchange organization. Members can also delegate role groups and management roles in the organization.

Note: If you select **Assign Organization Management** rights, the Microsoft application agent adds the user to the Organization Management group. The tool does not create the Dell EMC App Agent Exchange Admin Roles security group. If you do not select this option and do not select the **Skip Active Directory Authentication** option, the Microsoft Application Agent creates an Active Directory security group **Dell EMC App Agent Exchange Admin Roles** and adds the user to that group.



- Permissions that the Exchange Admin Configuration tool configures:
 - Security group memberships on the Microsoft application agent client host:
 - Local Administrator
 - Security group memberships on Domain Controller:
 - Remote Desktop Users
 - Exchange Security Group memberships:
 - Exchange Servers
 - Dell EMC App Agent Exchange Admin Roles, which include:
 - Exchange Roles

- o Database Copies
- o Databases
- o Disaster Recovery
- o Mailbox Import Export
- o Mail Recipient Creation
- o Mail Recipients
- o View-Only Configuration

Microsoft application agent for Exchange database backup

The Microsoft application agent uses block-based backup technology to back up Exchange Server databases in stand-alone and DAG environments. Block-based backup (BBB) technology tracks changed blocks of the Exchange database and log files as follows:

- A full backup backs up each selected Exchange database and log files.
- An incremental backup backs up only the changed blocks.

Block-based backups have reduced backup times because the backup process respectively backs up only the occupied disk blocks and changed disk blocks of the Exchange database and log files. During the backup, the application agent scans a volume or a disk where Exchange databases reside and backs up only changed blocks that are related to Exchange database.

Block-based backups provide instant access to the backups. These backups enable you to access databases using ItemPoint and perform granular restores.

Block-based backups use the following technologies:

- The Volume Shadow Copy Service (VSS) snapshot capability on Windows creates consistent copies of the source volume for backups.
- The Virtual Hard Disk (VHD), which is sparse, backs up data to the target device.

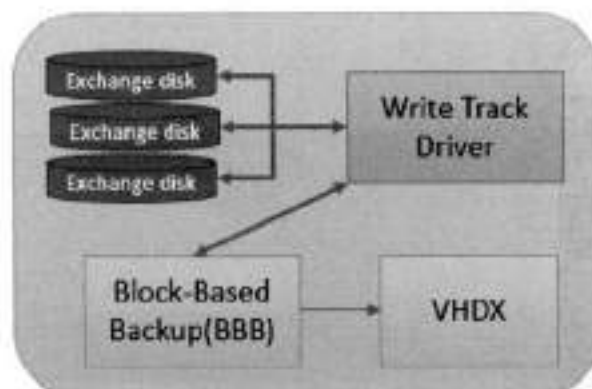


Figure 2. Block-based backups

Centralized backup

With centralized backups, the agent service coordinates the centralized backups as requested by the Data Manager. The backup is triggered according to the schedule, and Data Manager manages the complete protection life cycle.

The data protection attributes are specified when you create the centralized protection policy. These attributes include Type, Purpose, Assets, Schedule, Retention, and SLA. After you create the protection policy, the lockbox is automatically created, and the configuration information is saved in SQLite database called **configinfo.db**.

The following steps and Figure 3 describe the centralized backup workflow:

1. The backup schedule starts, and the ADM service triggers the protection policy which sends a REST API request to PowerProtect Data Manager agent service.
2. The PowerProtect Data Manager Agent service receives the requests and sends the backup request to the Microsoft application agent.
3. The Microsoft application agent requests Exchange VSS to generate snapshots for each database.
4. MsAgentsv.exe transfers the Exchange backup data blocks in VHDX format.
5. The data blocks are transferred from the snapshot to the PowerProtect DD series appliances. The first backup with the centralized policy is always full, and all incremental backups are virtual synthetic full.
6. Once the backup is completed, backup metadata information is stored in SQLite database.
7. The PowerProtect Agent service sends the backup-complete information to the ADM service, and the Data Manager UI reflects the backup completion status.

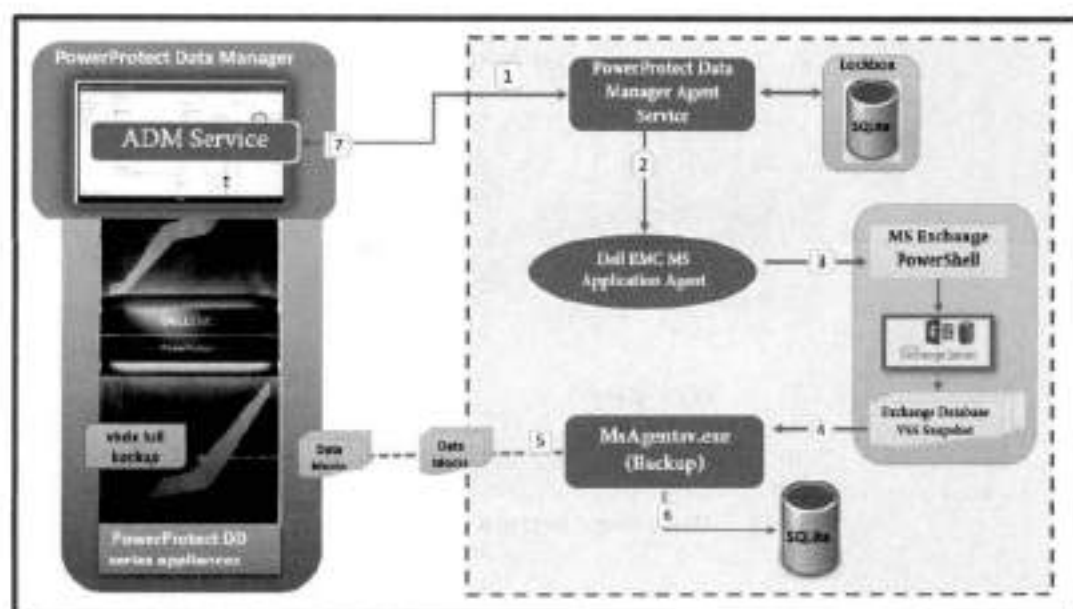


Figure 3. Centralized backup workflow

Federated backup of a DAG

A database availability group (DAG) environment can contain multiple passive copies of databases that are distributed across multiple Exchange servers. When you back up either active or passive database copies in the DAG environment, all DAGs use the federated backup method to best handle fail-over scenarios.

The federated backup method provides the following benefits:

- Allows backups of passive database copies to continue even when the passive database copies move among Exchange servers.
- Enables you to back up all DAG members, including stand-alone and public-folder-mailbox databases, by using a single save set. You are not required to perform a separate backup of each node.

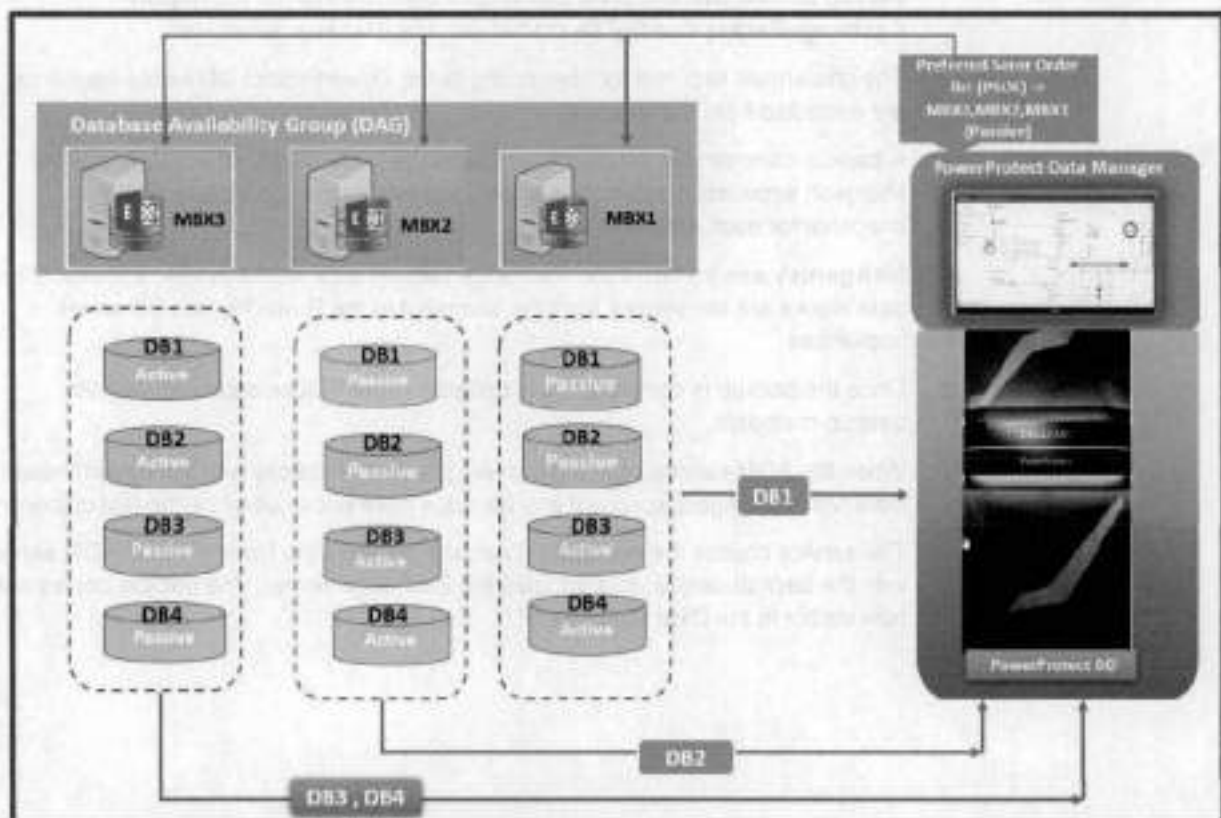


Figure 4. Federated backup of a DAG

The following steps describe the sequence of operation:

1. When the backup starts, the PowerProtect ADM service initiates the `getPreferredNode ()` PSOL.
2. The Exchange add-on fetches the preferred server order list of assets through PowerShell.

3. The PowerProtect Data Manager Backup Agent service triggers the preferred node list for the set of databases to be backed up.
4. The **MsAgentsv.exe** processes the backup request and indexes the backup metadata in SQLite.

Self-service backup

To enable self-service protection, when you create the Exchange protection policy, select **Self-Service Protection**. When performing a self-service stand-alone backup of a DAG asset, the backups appear under the DAG asset.

The Microsoft application agent supports full and incremental block-based backups.

Note: For self-service backups, do not select assets from multiple protection policies in the same backup request. This is a limitation of the Microsoft application agent.

The following steps and Figure 5 describe the workflow of a self-service backup:

1. To import the backup parameters to the object, the Exchange administrator or backup admins use Microsoft Exchange PowerShell to run the **Import-ExchangeBackupConfigFile** cmdlet with the **-Backup** parameter.
2. The credentials required for connecting to the PowerProtect DD series appliance are extracted from the lockbox.
3. A backup command is initiated from Exchange PowerShell, which launches the Microsoft application agent. The agent requests Exchange VSS to create a snapshot for each database.
4. **MsAgentsv.exe** transfers the Exchange backup data blocks in VHDX format. The data blocks are transferred from the snapshot to the PowerProtect DD series appliances.
5. Once the backup is completed, it is updated in the SQLite database with the backup metadata.
6. When the ADM service runs a discovery process, it checks with the PowerProtect Data Manager Agent service if any backups have occurred since the last discovery.
7. The service checks the metadata database, updates the PowerProtect ADM service with the backup details, and updates the Exchange server. The backup copies are now visible in the Data Manager UI.

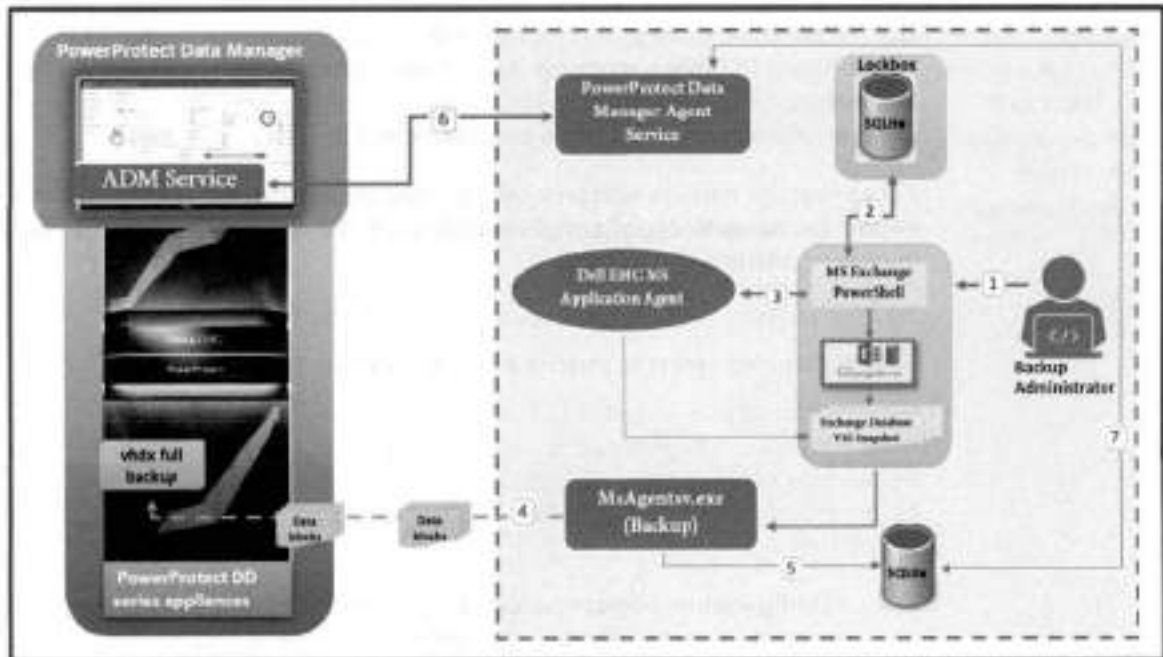


Figure 5. Self-service backup workflow

Parallelism for Microsoft application agent backup

To configure the parallelism setting for a centralized backup of Exchange databases, set the `clientParallelism` parameter value in the `userlockbox.cfg` file on the Exchange server host.

- The default parallelism setting for a centralized Exchange backup is **16**. You can override this default setting with the `clientParallelism` parameter setting in the configuration file.
- For a self-service Exchange backup, you can specify the parallelism with the `-Parallelism` parameter setting in the `Backup-Exchange` PowerShell cmdlet.

Based on the number of CPUs on the host and the parallelism setting, the application agent uses the following effective parallelism value for a centralized Exchange backup:

- **With 10 or more CPUs:** The effective parallelism equals the minimum number of CPUs and the parallelism setting, minus 4. For example, with 12 CPUs and a parallelism of 16, the effective parallelism value is 8 (12 minus 4).
- **With 4 to 9 CPUs:** The effective parallelism equals the minimum of the number of CPUs and the parallelism setting, minus 2. For example, with 8 CPUs and a parallelism of 10, the effective parallelism value is 6 (8 minus 2).
- **With fewer than 4 CPUs:** The effective parallelism equals the number of CPUs. For example, with 2 CPUs, the effective parallelism value is 2.

If the effective parallelism value is 8, then 8 threads are created for the Exchange backup, with each thread assigned to transfer either an EDB file or the related log files. An asset transfer is complete only when an EDB file and the related log files are both copied to the DD system. Using 8 threads, 4 assets are backed up in parallel from the host.

Back up Exchange server with Microsoft application agent PowerShell backup cmdlet

Use the Backup-Exchange PowerShell cmdlet to back up Exchange Server to a PowerProtect DD series appliance. All cmdlets support the standard common PowerShell parameters. The Microsoft article [Exchange Server PowerShell \(Exchange Management Shell\)](#) provides the list of common parameters and their description.

For self-service backups with protection policies created through Data Manager, run the **Import-ExchangeBackupConfigFile** cmdlet with the **-Backup** parameter to import the backup parameters to the object.

Syntax to perform stand-alone server backups

See the following syntax to perform stand-alone server backups:

```
[<configuration_object>] | Backup-Exchange -  
BackupViaBlockBasedBackup -ClientName <FQDN_of_Exchange_Server> -  
DataDomainHost <Data_Domain_hostname> -DataDomainHostPath  
/<Data_Domain_storage_path> -DataDomainUser <Data_Domain_username>  
[<optional_parameters>]
```

- **<configuration_object>** (optional): Specifies the configuration object that was imported using the Import-ExchangeBackupConfigFile cmdlet.
- **-ClientName <FQDN_of_Exchange_Server>**: Specifies the FQDN of the Exchange Server to use for indexing the backup.
- **-BackupViaBlockBasedBackup**: Specifies that the backup is a block-based backup. You can use the **-BBB** alias for the **-BackupViaBlockBasedBackup** parameter.
- **-DataDomainHost <Data_Domain_hostname>**: Specifies the PowerProtect DD series appliance server hostname. You can use the **-S**, **-SH**, **-DDHost**, or **-StorageHost** alias for the **-DataDomainHost** parameter.
- **-DataDomainHostPath /<Data_Domain_storage_path>**: Specifies the full path to the Data Domain storage unit for the backup. The PowerProtect DD series appliance user must have appropriate access rights to this path. You can use the **-Path**, **-DevicePath**, **-StoragePath**, **-StorageHostPath**, or **-DataDomainPath** alias for the **-DataDomainHostPath** parameter.
- **-DataDomainUser <Data_Domain_username>**: Specifies the PowerProtect DD series appliance username. Full credentials are retrieved from the lockbox to authenticate with the host. **-DDUser**, **-StorageUser** You can use the **-DDUser** or **-Storage User** alias for the **-DataDomainUser** parameter.

The following is an example of a backup command:

```
Backup-Exchange -ClientName myexchange.msapp.com -  
BackupViaBlockBasedBackup -DataDomainHost myDD.1ss.example.com -  
DataDomainPath /SU_DD163 -DataDomainUser DD163_user
```

The following is an example of a backup command with a configuration object:

```
$serverinfo | Backup-Exchange
```

Syntax to perform federated backup

See the following syntax to perform a federated backup:

```
[<configuration_object>] | Backup-Exchange -
BackupViaBlockBasedBackup -<ClientName
<FQDN_of_Exchange_Server_DAG> -DataDomainHost
<Data_Domain_hostname> -DataDomainHostPath
/<Data_Domain_storage_path> -DataDomainUser <Data_Domain_username>
[[-BackupActive] | [-BackupPassive] | [-BackupPreferred]] [-
IncludeStandaloneDatabases] [-
ServerOrderList<comma_separated_list_of_servers>]
[<optional_parameters>]
```

- **-ClientName <FQDN_of_Exchange_Server_DAG>**: Specifies the FQDN of the database availability group instance to use for indexing the backup.
- **{-BackupActive | -BackupPassive | -BackupPreferred}** (optional): Specifies that the database backup preference is either active (-BackupActive), passive (-BackupPassive), or preferred (-BackupPreferred).
- **-IncludeStandaloneDatabases** (optional): Specifies to include stand-alone databases and public folder databases in the backup.
- **-ServerOrderList <comma_separated_list_of_servers>** (optional): Specifies the preferred Exchange Server order list if you must select multiple copies. Separate multiple servers with commas.

The following is an example is of a federated backup command:

```
Backup-Exchange -Identity TestDB,'Mailbox Database 1250665181' -
ClientName DAG1.msapp.com -BackupViaBlockBasedBackup -
DataDomainHost myDD.lss.example.com -DataDomainPath /SU_DD163 -
DataDomainUser DD163_user -Preferred -ServerOrderList node1, node2
-IncludeStandaloneDatabases
```

The following is an example of a federated backup command with a configuration object:

```
$serverinfo | Backup-Exchange -Identity TestDB,'Mailbox Database
1250665181'
```

Optional parameters for the Backup-Exchange cmdlet

The following list describes the optional parameters for the Backup-Exchange cmdlet:

- **-Incremental**: Specifies that the backup level is a block-based incremental backup. If you do not specify this parameter, the backup is taken at the full level.
- **-Retention +<number>{d | m | w | y}**: Specifies the period in which to retain a backup. After the period passes, the backup expires. The default retention period is 30 days. The maximum retention date is 02-07-2106.
- **-Identity <database identity>**: Specifies the identity of the database to back up. If you do not specify this parameter, the operation backs up all databases.
- **-LockBoxPath <full_path_to_lockbox>**: Specifies the folder that contains the lockbox file, which contains encrypted information about the registered hosts and

the corresponding usernames in pairs. Each pair is associated with a password that the backups use.

- **-ExeFileName <msagentsv.exe_path>**: Specifies the full path to the application program executable msagentsv.exe. Use this option only for diagnosis. In normal operations, the cmdlet automatically locates the installed application.
- **-AsJob (\$true | \$false)**: Runs the cmdlet as a background job. The command returns an object that represents the job and displays the command prompt. You can continue to work in the session during the job.
- **-Parallelism <parallelism_value>**: Specifies the parallelism setting for the backup.

Note: For a complete list of available options, see the [PowerProtect Exchange Server Guide](#).

Restoring Exchange Server databases

Introduction

You can perform database restores or granular-level restores directly to the Exchange application host using the Microsoft application agent. The agent supports the following types of database restores:

- **Normal restore:** Restore of a database to the original source database.
- **Alternate database restores:** Restore of a database to another database that is different from the source database.

Prerequisite for Exchange restore operation

You must run the **set-mailboxdatabase** cmdlet to allow an Exchange database to be restored from a backup.

```
set-mailboxdatabase <mailbox_database> -AllowFileRestore $true
```

- **<mailbox_database>**: Specifies the name of the database that is the target for the restore operation.
- **-AllowFileRestore \$true**: Specifies to allow restore operations for the database.

Note: Run this command for each target database for the restore operation.

Restore a backup to the source database

Use the **Restore-Exchange** cmdlet with the following syntax to restore a database to the source location (normal restore):

```
[<configuration_object>] | Restore-Exchange -NormalRestore (-BackupID <backup_ID> [-Identity <identity>] | -Backup <backup_object>) -ClientName <FQDN_of_Exchange_Server> -DataDomainHost <Data_Domain_hostname> -DataDomainHostPath /<Data_Domain_storage_path> -DataDomainUser <Data_Domain_username> <optional_parameters>
```

- **<configuration_object>** (optional): Specifies the configuration object that was imported using the **Import-ExchangeBackupConfigFile** cmdlet.

Note: For Data Manager centralized and self-service workflows, run the **Import-ExchangeBackupConfigFile** cmdlet with the **-Restore** parameter to import the configuration parameters to the object.

- **-NormalRestore:** Specifies that the database is being restored to the original source location. You can use the **-Restore** alias for the **-NormalRestore** parameter.
- **{-BackupID <backup_ID> [-Identity <identity>]} | -Backup <backup_object>:** Specifies the backup to restore using either the backup identity or object. You must specify only one of the following options:
 - **-BackupID <backup_ID>:** Use a backup ID. Optionally, specify **-Identity <database_ID>** with **-BackupID** to specify the identity of one or more databases to restore.
 - **-Backup <backup_object>:** Use a backup object. You can retrieve the backup ID and object from the **Backup-Exchange** or **Get-ExchangeBackup** cmdlet output. The following example restores the database TestDB by using a backup ID.

```
Restore-Exchange -NormalRestore -BackupID nsapp_bbb:
1458138556 -Identity TestDB -ClientName myDB.msapp.com -
DataDomainHost ledmd035.lss.example.com -DataDomainHostPath
/SU_DD163 -DataDomainUser DD163_user
```

Restore a backup to an alternate database

Note: Before you perform a copy or alternate-database restore, ensure that the target database exists. Use the **Restore-Exchange** cmdlet with the following syntax to restore a database to an alternate location (copy restore):

```
[<configuration_object>] | Restore-Exchange -CopyRestore -BackupID
<backup_ID> -Identity <identity> -RestoreDatabaseIdentity
<target_identity> -ClientName <FQDN_of_Exchange_Server> -
DataDomainHost <Data_Domain_hostname> -DataDomainHostPath
/<Data_Domain_storage_path> -DataDomainUser <Data_Domain_username>
[<optional_parameters>]
```

- **-CopyRestore:** Specifies that the database is being restored to an alternate location. You can use the **-Alternate** alias for the **-CopyRestore** parameter.
- **-BackupID <backup_ID>:** Specifies the backup ID to restore. You can retrieve the backup ID from the **Backup-Exchange** or **Get-ExchangeBackup** cmdlet output.
- **-Identity <database_ID>:** Specifies the identity of one or more databases to restore.
- **-RestoreDatabaseIdentity <target_identity>:** Specifies the target identity of the alternate database to restore to. You can use the **-RestoreDB**, **-Target**, **-RDB**, or **-RestoreDatabaseID** alias for the **-CopyRestore** parameter.

The following shows an example of restoring the database TestDB to an alternate database (AlternateDB) using a backup ID:

```
Restore-Exchange -CopyRestore -BackupID msapp_bbb: 1458138556 -
Identity TestDB -RestoreDatabaseIdentity AlternateDB -ClientName
ledmf175.msapp.com -DataDomainHost myDD.lss.example.com -
DataDomainHostPath /SU_DD163 -DataDomainUser DD163_user
```

Granular-level restores

To recover granular-level Exchange Server data, you must first mount the backup using the **Mount-ExchangeBackup** PowerShell cmdlet. When the backup is mounted, you can browse and recover granular items, such as mailboxes or folders, with ItemPoint for Microsoft Exchange Server.

To perform granular-level restores, first mount the backups. Use the **Mount-ExchangeBackup** cmdlet with the following syntax to mount the backups:

```
[<mount_object> = <configuration_object>] Mount-ExchangeBackup [-
BackupID <backup_ID> [-Identity <identity>] | -Backup
<backup_object>] -ClientName <FQDN_of_Exchange_Server> -
DataDomainHost <Data_Domain_hostname> -DataDomainHostPath
/<Data_Domain_storage_path> -DataDomainUser <Data_Domain_username>
[<optional_parameters>]
```

Result: The backup is mounted in a path like the following:

```
C:\Program
Files\UPSAPPS\MSAPPAGENT\tmp\BBBMountPoint\131248297060279537_(4A6
0AF18-B6ED-4BBD1C9-2618F1AC1041)_5832\Program
Files\Microsoft\Exchange Server\VI5\Mailbox\DB2\
```

The mounted items are unmounted after you restart the host.

The following are examples of the Mount-ExchangeBackup cmdlet:

Mount all databases of a backup using a backup object:

```
Mount-ExchangeBackup -Backup $backups [0] -ClientName
ledmf175.msapp.com -DataDomainHost myDD.lss.example.com -
DataDomainHostPath /SU_DD163 -DataDomainUser DD163_user
```

Mount a single mailbox database database3 using a backup object and identity:

```
$mount = $serverinfo | Mount-ExchangeBackup -Backup $backup[0] -
Identity database3
```

Browse and recover granular-level data with ItemPoint for Microsoft Exchange Server

The ItemPoint for Exchange Server User Guide provides more information about performing granular-level recovery of Exchange data. Perform the following steps:

1. Launch ItemPoint.
2. In ItemPoint, launch the Restore wizard.

- On the **Source Selection** page, select the source and specify the EDB and log file path from the mounted volume that contains the Exchange backup data (see the following screen). Click **Next**.

- On the **Target Selection** page, click **Skip**.

- Follow the Data Wizard prompts to complete the granular-level recovery.
- Once the granular-level recovery is complete, dismount the backup.

Cloud tiering and replication

During the protection policy creation, you can add the replication to a remote PowerProtect DD series appliance as the replication target.

The Data Manager cloud tier feature works in tandem with the Data Domain Cloud Tier feature to move Data Manager backups from the PowerProtect DD series appliance to the cloud. This feature provides long-term storage of Data Manager backups by seamlessly and securely tiering data to the cloud. From the Data Manager UI, you configure the cloud tier to move Data Manager backups from PowerProtect DD series appliance to the cloud. You can also perform seamless recovery of these backups. Data Domain cloud storage units must be preconfigured on the PowerProtect DD series appliance before they are configured for cloud tier in the Data Manager UI. See the [Data Domain Operating System Administration Guide](#) for more information.

Both Exchange centralized and self-service protection policies support cloud tiering. You can create the cloud tier schedule from both primary and replication stages. Schedules must have a minimum weekly recurrence and a retention time of 14 days.

Disaster recovery of Exchange Server

Introduction

When a disaster scenario occurs, the Microsoft application agent can provide disaster recovery of data that is on both a PowerProtect DD series appliance server and Data Domain Cloud Tier.

Follow these steps to perform a disaster recovery on the new disaster recovery host.

1. Start the Exchange Server application and the required services.
2. Create the databases that existed before the disaster and ensure that the databases are in the mounted state.
3. Perform a restore of the databases.

Perform disaster recovery from the Data Domain Cloud Tier

The Microsoft application agent provides a command-line tool to complete disaster recovery of save sets that are in a Data Domain Cloud Tier. After an MTree is recovered according to the disaster recovery procedure, you must restore the backup indexes from the Data Domain Cloud Tier.

When the Microsoft application agent moves a backup to the cloud, the index files are maintained on the active tier. A copy of the index files is created and moved to the cloud tier for long-term retention.

After an MTree is restored during a disaster recovery, all the files that resided only on the active tier are lost and unavailable. Only the files that were moved to the cloud are available. In this case, you must run `msagentadmin` administration with the `--dr-recall` parameter to restore the indexes.

After the indexes are recalled to the active tier, the data save sets for the same time range are also recalled unless you type `n` when prompted and browse to the recall of the found save sets `[y/n]`. If you choose to not recall the save sets, you can manually recall the save sets later.

Use the `msagentadmin` administration command with the following syntax to recall the indexes to the active tier:

```
msagentadmin.exe administration --dr-recall --ddhost
"<Data_Domain_server_name>" --ddpath
"<name_and_path_of_storage_unit>" --dduser "<DDBoost_username>" --
appID <application_ID>
```

- `--dr-recall`: Specifies an operation to recall save sets for disaster recovery. You can use the `-M` alias for the `--dr-recall` parameter.
- `--ddhost "<name>"`: Specifies the name of the PowerProtect DD series appliance server that contains the storage unit, to which you backed up the databases.
- `--ddpath "/<storage_unit_name_and_path>"`: Specifies the name and the path of the storage unit, to which you backed up the databases.
- `--appID "<application_ID>"`: Specifies the application ID (namespace) to locate backups. Specify `msapp_bbb` for Exchange Server. You can use the `-n` alias for the `--appID` parameter.

Consider the following example commands to perform disaster recovery of Exchange Server with data on a PowerProtect DD Series Appliance Cloud.

Tier device:

- Cloud tier disaster recovery recall command without a configuration file:

```
msagentadmin administration --dr-recall --tier --after
1481104962 --before 1481105533 -appID msapp_bbb --ddhost
"10.70.102.111" --ddpath "/mt1" --dduser "ost" --confirm -
client myDD.msapp.com --debug 9
```

- Cloud tier disaster recovery recall command with a configuration file:

```
msagentadmin.exe administration --dr-recall --tier --after
1481104962 --before 1481105533 --appID msapp_bbb --confirm --
config c:\temp\config_pp.txt --debug 9
```

Conclusion

Dell EMC PowerProtect Data Manager enables complete control of Microsoft Exchange database backup and disaster recovery to backup administrators. The advanced integration between PowerProtect Data Manager and Microsoft Exchange provides a fast and efficient database backup and restore solution.

Technical support and resources

Dell.com/support is focused on meeting customer needs with proven services and support.

[Storage and data protection technical white papers and videos](#) provide expertise that helps to ensure customer success with Dell EMC storage and data protection products.

For more information, see the following related resources:

- [PowerProtect Data Manager Microsoft Application Agent Exchange Server User Guide](#)
- [PowerProtect DD Series Appliance Operating System Administration Guide](#)
- [PowerProtect Data Manager Administration and User Guide](#)

Dell PowerProtect Data Manager: Deployment Best Practices

December 2022

H18564.4

White Paper

Abstract

This white paper explains PowerProtect Data Manager deployment best practices and includes deployment and setup requirements for a new PowerProtect Data Manager.

Dell Technologies

Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2019-2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA December 2022 H18564.4.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Contents

Executive summary	4
PowerProtect Data Manager overview	5
PowerProtect Data Manager deployment methods	6
PowerProtect Data Manager deployment considerations	7
Deployment of PowerProtect Data Manager	15
Post-installation steps for PowerProtect Data Manager	16
Multiple VLAN configurations	32
Scalability limits for PowerProtect Data Manager	37
Conclusion	38
Technical support and resources	38

Executive summary

Overview

Data protection has become an integral and essential part of any successful business. The need to provide a powerful, scalable, and simple disaster and operational recovery solution is at an all-time high. IT teams are also looking for a solution that is scalable, easy to implement, efficient to use and handles the workload of their small and medium size environments. To meet midmarket industry demands, Dell Technologies offers Dell PowerProtect Data Manager.

PowerProtect Data Manager enables the transformation from traditional, centralized protection to a Software-as-a-Service (SaaS) model based on a self-service design. This design ensures that you can enforce compliance and other business rules even when backup responsibilities are decentralized to individual database or application administrators.

Some key differentiators for Data Manager are:

- Software-defined backup appliance with integrated deduplication for data protection, replication, and reuse
- Self-service for data owners concerned with central IT governance
- SaaS-based management, compliance, and predictive analytics
- Multidimensional with scale-up and scale-out flexibility and all flash performance
- Microservices architecture for ease of deployment, scaling and upgrading
- Multicloud that is optimized with integrated cloud tiering and cloud disaster recovery

This paper focuses on the PowerProtect Data Manager deployment requirements and best practices.

Audience

This white paper is intended for customers, partners, and employees who want to better understand, evaluate, and explore deployment requirements and best practices of PowerProtect Data Manager. Familiarity with PowerProtect DD series appliances is required.

Revisions

Date	Description
July 2019	Initial release
October 2020	Revision
July 2021	Revision
July 2022	Content update
December 2022	Updates to hyperlinks and scalability limits; template and editorial updates

We value your feedback

Dell Technologies and the authors of this paper welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

Author: Vino Jeyakanth

Note: For links to other documentation for this topic, see the [PowerProtect Data Manager Info Hub](#).

PowerProtect Data Manager overview

PowerProtect Data Manager integrates multiple data protection products within the Dell Data Protection portfolio to enable data protection as a service. PowerProtect Data Manager enables new data paths with provisioning, automation, and scheduling that enable a data protection team to embed protection engines into their infrastructure for high-performance backup and recovery.



Figure 1. PowerProtect Data Manager overview

PowerProtect Data Manager offers the following benefits:

- Uses an agent-based approach to discover the protected and unprotected databases on an application server.
- Enables governed self-service and centralized protection by:
 - Monitoring and enforcing Service Level Objectives (SLOs)
 - Identifying violations of Recovery Point Objectives (RPO)
 - Applying retention locks on backups for all asset types
- Supports deploying an external VM Direct appliance to move data with the VM Direct Engine.
- Supports the vRealize Automation data protection extension, which enables provisioning of virtual machines with Data Manager protection, manual backup, and restore to the original or a new location.

- Supports integration of Cloud Disaster Recovery (Cloud DR), including workflows for Cloud DR deployment, protection, and recovery operations in the AWS or Azure cloud.
- 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:
 - VMware virtual machines
 - File systems
 - VMAX storage groups
 - Kubernetes clusters
 - Microsoft Exchange and SQL databases
 - Oracle databases
 - SAP HANA databases
 - Network-attached storage (NAS) shares
- For details about version support and compatibility, see the [PowerProtect Data Manager Compatibility Matrix](#).
- Supports PowerProtect Search, which enables backup administrators to quickly search for and restore VM file copies. The Search Service can be enabled by adding a search node to the configurable Search Engine that is auto deployed during the Data Manager installation.
- Provides a RESTful interface that allows the user to monitor, configure, and orchestrate Data Manager. Customers can use the APIs to integrate their own automation framework or quickly write new scripts with the help of easy-to-follow tutorials.

PowerProtect Data Manager deployment methods

Introduction to deployment methods

You can deploy PowerProtect Data Manager using an Open Virtualization Appliance (OVA) or a machine image. Each method has its own considerations for the deployment itself and the functionality of PowerProtect Data Manager after its deployment.

OVA deployments

Considerations of OVA deployments include the following:

- PowerProtect Data Manager can be deployed to on-premises virtual hosts or to cloud-based environments that include VMware Cloud on Dell, VMware Cloud (VMC) on Amazon Web Services (AWS), and Azure VMware Solution (AVS) on Microsoft Azure.
- OVAs are deployed using the vSphere Client.
- Deployed PowerProtect Data Manager instances do not detect their environment. The environment must be manually selected during the deployment process for the instances to be appropriately configured.

- PowerProtect Data Manager and DDVE cannot be deployed simultaneously from the same interface.

Machine-image deployments

For machine-image deployments, consider the following information:

- PowerProtect Data Manager can only be deployed to virtual hosts on Amazon Web Services (AWS), Microsoft Azure, or Google Cloud Product (GCP). These virtual hosts cannot be in an environment that includes VMware Cloud (VMC) or Azure VMware Solution (AVS), although any deployed PowerProtect Data Manager can still protect resources in those environments.
- Machine images are deployed using the web-based user interface of a cloud provider.
- Deployed PowerProtect Data Manager instances detect their environment and are automatically appropriately configured.
- PowerProtect Data Manager and DDVE can be deployed simultaneously and from the same interface.

This white paper describes how to deploy PowerProtect Data Manager using an OVA. For information about how to deploy PowerProtect Data Manager using a machine image, see the following guides:

- [PowerProtect Data Manager Amazon Web Services Deployment Guide](#)
- [PowerProtect Data Manager Azure Deployment Guide](#)
- [PowerProtect Data Manager Google Cloud Platform Deployment Guide](#)

PowerProtect Data Manager deployment considerations

Deployment introduction

Planning the environment for deploying PowerProtect Data Manager using OVA plays an important prerequisite function. You must plan for adequate resources to achieve optimal performance of PowerProtect Data Manager.

Planning VMware vCenter resources

PowerProtect Data Manager Appliance

The following table outlines the minimum resource requirements to deploy a PowerProtect Data Manager OVA in VMware vSphere 6.0 and later:

Table 1. Resource requirements for PowerProtect Data Manager OVA deployment in vSphere 6.0 and later

Specification	Value
CPU	10 CPU cores
Memory	24 GB RAM
Disk	<ul style="list-style-type: none"> • Disk 1: 100 GB • Disk 2: 500 GB • Disk 3 and 4: 10 GB each • Disk 5–7: 7 GB each

Specification	Value
Virtual Disk Format	Thick provision lazy zeroed
Network interface card (NIC)	1 GB
Internet Protocol	IPv4 only
For application-aware backup on VM	<ul style="list-style-type: none"> vCenter version 6.5 or later VMware ESXi server version 6.5 or later VMware tool version 10.1 or later
For Cloud DR	<ul style="list-style-type: none"> 14 CPU cores (10 for PowerProtect Data Manager and 4 for Cloud DR) 28 GB RAM (24 GB for PowerProtect Data Manager and 4 GB for Cloud DR)

PowerProtect Search Engine

The PowerProtect Search Engine enables backup administrators to quickly search and restore VM file copies. The Search Service can be enabled by adding a search node to the configurable Search Engine that is auto deployed during PowerProtect Data Manager installation.

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 search cluster. Adding a node enables the indexing feature.

Each search engine node must meet the following system requirements:

Table 2. System requirements for search engine node

Specification	Value
CPU	4 vCPU * 2 GHz (4 virtual sockets, 1 core for each socket)
Memory	8 GB RAM
Disk	<ul style="list-style-type: none"> Three disks - 50 GB each One disk - 1 TB
Internet Protocol	IPv4 only
Network interface card (NIC)	One vmxnet3 NIC with one port

Note: You can add up to a maximum of five search engine nodes for a single PowerProtect Data Manager. One search node can index maximum 1 billion files or 1,000 virtual machines.

PowerProtect VM Direct protection engine

PowerProtect Data Manager comes prebundled with an embedded VM Direct Engine. The engine is automatically used as a fallback proxy for performing backup and restore operations when the added external proxies fail or are disabled. The VM Direct Engine facilitates data movement for both virtual machine protection policies and Kubernetes cluster protection policies.

Dell Technologies recommends that you always deploy an external protection engine, also known as a VM proxy, because the embedded proxy has limited capacity for performing parallel backups. The following table details the requirements for the external VM Direct Engine:

Table 3. External VM Direct Engine requirements

Specification	Value
CPU	4 vCPU * 2 GHz (4 virtual sockets, 1 core for each socket)
Memory	8 GB RAM
Disk	Disk 1: 59 GB Disk 2: 98 GB
Internet Protocol	IPv4 only
Network interface card (NIC)	One vmxnet3 NIC with one port
SCSI controller	4 (maximum)

Notes:

- Total number of external VM Direct Engines supported with a single vCenter server is 25, although the recommended number is 7.
- Network settings such as Gateway, IP Address, Netmask, and Primary DNS are important to specify.
- Each external VM Direct Engine can manage a maximum of 25 VM backup and recovery sessions.
- The embedded VM Direct Engine supports four backup and restore sessions.

Best Practices:



- Create a dedicated PowerProtect vCenter user, and avoid using the vCenter administrator
- Install VMware Tools on each virtual machine
- Use Hotadd transport mode for faster backups and restores and less exposure to network routing, firewall, and SSL certificate issues
- Avoid deploying virtual machines with IDE virtual disks

Planning the protection storage

PowerProtect Data Manager integrates multiple data protection products within the Dell Data Protection portfolio to enable data protection as a service. It enables new data paths with provisioning, automation, and scheduling that allow a data protection team to embed protection engines into the infrastructure for high-performance backup and recovery.

The PowerProtect Data Manager UI enables users with administrator credentials to add the following storage types:

- PowerProtect DD Management Center
- External PowerProtect DD series appliance

Note: You can also add a DD system in High Availability (HA) mode.

The following table shows the supported versions of DD series appliances:

Table 4. Supported versions of DD series appliances

PowerProtect DD series appliance	Supported versions
Hardware	DD990, DD4500, DD7200, DD9500, DD6300, DD6800, DD9300, DD9800, DD3300
Operating system	DDOS 6.1.2 or higher
PowerProtect DD Management Center (DDMC)	DDMC 6.1.0.x, 6.1.x with DDOS 6.1.x and higher
DD Virtual Edition (DDVE)	DDVE 4.0 and above



Best Practice: Create a dedicated PowerProtect Data Domain BOCST user and avoid using sysadmin account for Data Domain discovery.

Note: When a PowerProtect DD Management Center is added, PowerProtect Data Manager discovers all the supported DD series appliances that are managed by the PowerProtect DD Management Center.

Planning networking

The following sections outline the networking requirements for deploying PowerProtect Data Manager.

IP and DNS requirement

Networking requirements are as follows:

- Unique IP addresses must be allocated to the PowerProtect Data Manager, VM Search Engine, and VM Direct Engine. Only IPV4 IP addresses are supported.
- NTP servers are recommended to sync PowerProtect Data Manager with NTP server.
- DNS server and default gateway servers should also be specified during install. You can configure up to three DNSs.
- Forward and reverse DNS lookups are recommended.
- When configuring PowerProtect Data Manager, do not use an IP address in the 172.24.0.192 /26 subnet. IP addresses from 172.24.0.192 through 172.24.0.255 are reserved for the private Docker network.
- vCenter registration and proxy deployment fail if the PowerProtect Data Manager server is deployed in the same private network as the internal Docker network.

Firewall and port requirements

PowerProtect Data Manager is a single node in a virtual appliance that uses the Linux SLES 12 firewall to protect and limit external access to the system. PowerProtect Data Manager uses a direct socket connection to communicate and move data internally and across the network to the required service with minimal overhead.

To enable communication between the PowerProtect Data Manager system and other applications, PowerProtect Data Manager configures firewall rules for ports that are used for inbound and outbound communication. The following table shows the port requirements for PowerProtect Data Manager:

Table 5. PowerProtect Data Manager port requirements

Description	Communication	Port
SSH communications	Bi-directional communication between the SSH client and the PowerProtect Data Manager appliance	22 TCP/UDP
SQL, Oracle, Exchange, SAP HANA, file system	Bi-directional communication between the PowerProtect Data Manager agent and the PowerProtect Data Manager appliance Requirement applies to Application Direct and VM Direct.	7000 TCP
REST server	Bi-directional communication between the HTTP client and the PowerProtect Data Manager appliance	8443 TCP
RESTAPI server – VM Direct	Bi-directional communication between the PowerProtect Data Manager agent and the PowerProtect Data Manager appliance Requirement applies to SQL VM application aware.	8443 TCP
UI redirect	Inbound only	<ul style="list-style-type: none"> • 80 TCP • 443
LDAP	Outbound only	<ul style="list-style-type: none"> • 389 TCP/UDP • 636 TCP
Discovery (devices)	Outbound between the PowerProtect Data Manager appliance and the device	<ul style="list-style-type: none"> • 3009 TCP—Storage Direct and DD system • 5989 TCP—SMI-S • 443 TCP—XtremIO • 7225 TCP—RecoverPoint
PowerProtect Data Manager agent	Bi-directional communication between the database hosts and the PowerProtect Data Manager appliance This requirement applies to both Application Direct and VM Direct.	7000 TCP
Embedded VM Direct service	Outbound	9090 TCP

Description	Communication	Port
PowerProtect controller	Outbound between the PowerProtect Data Manager appliance and PowerProtect Controller on the Kubernetes cluster PowerProtect Data Manager uses this port to pull the logs from the controller pod.	30095 TCP
PowerProtect DD series appliance	Bi-directional port should be open between DD series appliance and External VM Direct or application hosts.	<ul style="list-style-type: none"> • 111 TCP • 2049 TCP • 2052 TCP
vCenter	Bi-directional between the PowerProtect Data Manager and vCenter for discovery, initiating Hot Add transport mode, restores including instant access restore.	<ul style="list-style-type: none"> • 443 HTTPS • 7444 TCP

Note: To get a detailed list, see the [PowerProtect Data Manager Security Configuration Guide](#).



Best Practices:

- Verify all components have network connectivity to each other
- Configure forward and reverse lookup addresses

Planning application hosts

Dell Technologies recommends preinstalling the supported application agents:

- PowerProtect Data Manager Oracle RMAN Agent
- PowerProtect Data Manager Microsoft Exchange server Agent
- PowerProtect Data Manager Microsoft SQL server Agent
- PowerProtect Data Manager File System Agent
- PowerProtect Data Manager SAP HANA Agent

User permission requirements for supported platforms

The following table shows an overview of the permissions requirements for all supported platforms.

Note: Check individual guides for more details about permissions for each workload.

Table 6. User permissions requirements for supported platforms

Application workload	User permissions requirements				Documentation
	Installation and discovery	Backup: Self-service	Backup: Centralized	Recovery	
SQL – Application Direct	Any user with local administrator privileges	Any user with local administrator privileges	Any user with local administrator privileges	Any user with local administrator privileges	PowerProtect Data Manager Microsoft Application Agent SQL Server User Guide

SQL – Application Aware	Any user with local administrator privileges	Not applicable	Any user with local administrator privileges	Any user with local administrator privileges	PowerProtect Data Manager Microsoft Application Agent SQL Server User Guide
Microsoft Exchange	Any user with local administrator privileges Note: To create a user using the App Agent Exchange Admin Configuration tool, log in with domain administrator permissions.	Exchange User configured using the App Agent Exchange Admin Configuration tool	Exchange User configured using the App Agent Exchange Admin Configuration tool	Exchange User configured using the App Agent Exchange Admin Configuration tool OR Any specific user group privileges to the system account and user account that will perform the restore operations	PowerProtect Data Manager Microsoft Application Agent Exchange Server User Guide
SAP HANA	Operating system root user	<ul style="list-style-type: none"> System database: System database user with backup administrator, catalog read role Tenant database: Either system database user with database administrator role or tenant database administrator with backup administrator, catalog read roles 	<ul style="list-style-type: none"> HANA studio: System database: System database user with backup administrator, catalog read role Tenant database: Either system database user with database administrator role or tenant database administrator with backup administrator, catalog read roles HANA CLI Hana (<sid>adm) user from CLI 	Database administrator or Hana <sid>adm operating system user	PowerProtect Data Manager SAP HANA Agent User Guide

PowerProtect Data Manager deployment considerations

Oracle	<p>Oracle user for RMAN agent install (.install.sh)</p> <p>And</p> <p>Operating system root user for agent service (PowerProtect Agent RPM)</p>	<p>Oracle operating system user (by default)</p> <p>OR</p> <p>Oracle database user (Oracle 12c and later sysbackup privilege is required of any database user) Precedence (high to low): Wallet/DB/OS Authentication</p>	<p>Oracle operating system user (by default)</p> <p>OR</p> <p>Oracle database user (Oracle 12c and later sysbackup privilege is required of any database user) Precedence (high to low): Wallet/DB/OS Authentication</p>	<p>Oracle operating system user (by default)</p> <p>OR</p> <p>Oracle database user (Oracle 12c and later sysbackup privilege is required of any database user) Precedence (high to low): Wallet/DB/OS Authentication</p>	<p>PowerProtect Data Manager Oracle RMAN Agent User Guide</p>
File system	<ul style="list-style-type: none"> For Linux: Operating system root user For Windows: Any user with local administrator privileges 	<ul style="list-style-type: none"> For Linux: Root user For Windows: Any user with local administrator privileges 	<ul style="list-style-type: none"> For Linux: Root user For Windows: Any user with local administrator privileges 	<ul style="list-style-type: none"> For Linux: Root user For Windows: Any user with local administrator privileges 	<p>PowerProtect Data Manager File System User Guide</p>
VMware	<p>A vCenter user account at the root level of the vCenter that is strictly dedicated for use with VM Direct protection engine</p> <p>Note: Avoid using vCenter administrator account.</p>	Not applicable	<p>Dedicated vCenter user account at the root level of vCenter</p> <p>Note: The User Guide provides a full list of user account privileges.</p>	<ul style="list-style-type: none"> Full VM restore: Dedicated vCenter user account at the root level of vCenter FLR restore: Any user with local administrator privileges 	<p>PowerProtect Data Manager Virtual Machine User Guide</p>
Storage Direct	<p>For Linux – root user</p> <p>For Windows- Any user with local administrator privileges</p>	<p>Ddboost and DdVdiskUser specified in the lockbox configuration file must have the administrator role.</p>	<p>Ddboost and DdVdiskUser specified in the lockbox configuration file must have the administrator role.</p>	<p>Ddboost and DdVdiskUser specified in the lockbox configuration file must have the administrator role.</p>	<p>PowerProtect Data Manager Storage Direct User Guide</p>

To provide local administrator privileges correctly to a domain user on Windows server:

1. On the Active Directory server, create a domain user account (dns\domain user) or use an existing domain user.
2. Make the user a member of "Backup Operations" and "Remote Desktop Users."
3. Log in to the application host as system administrator.
4. Go to **Control Panel > User Accounts > Manage User Accounts**, and add the new user with local administrator privileges.
5. Click **Start > Administrative Tools > Local Security Policy > User Rights Assignment > Log on as a Service**, and add the new user.
6. Disable UAC.

This domain user account can now be used to perform SQL, Exchange, and file system application agent installations.

Planning licensing

The available license types are as follows.

- **Trial:** Applied automatically on installation of PowerProtect Data Manager and enabling full use of the product for up to 90 days without applying a license key. When the trial period ends, PowerProtect Data Manager continues to operate with full functionality so that you can apply a permanent license.
- **Front-end protected capacity by terabyte (FETB):** The primary model of e-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.

Note: When upgrading from a previous release, any existing license, and its associated Secure Remote Services connections (SupportAssist replaces Secure Remote Services in PowerProtect Data Manager 19.8.) are removed from the system and replaced with a 90-day trial license. Therefore, a valid FETB license for PowerProtect Data Manager and any associated Secure Remote Services connections must be reinstalled.

To obtain the XML license file from the Dell license management website, you must have the License Authorization Code (LAC), which is emailed from Dell. If the LAC is misplaced or not received, contact a Dell technical support representative.

Deployment of PowerProtect Data Manager

PowerProtect Data Manager Appliance is easy to install and configure. The PowerProtect Data Manager Open Virtual Appliance (OVA) can be deployed using one of the following methods:

- **Manually deploying the OVA to a vCenter server—**Use this method to deploy the OVA to a stand-alone or cluster host, while logged into the vCenter server. Configuration of the network settings is supported during the deployment.

- Manually deploying the OVA to an ESXi host—Use this method to deploy the OVA while logged in to an ESXi host. Use the VM console to configure the network settings after the deployment completes.

Note: Enter the network details correctly in the network section of the OVA deployment flow; otherwise, the appliance will not connect after deployment.

Once the PowerProtect Data Manager OVA deployment is completed, all the services and components of the software are accessible.

Deploy OVF Template

✓ 1 Select an OVF template
✓ 2 Select a name and folder
✓ 3 Select a compute resource
✓ 4 Review details
✓ 5 Configuration
✓ 6 Select storage
✓ 7 Select networks
✓ 8 Customize template
➔ **Ready to complete**

Ready to complete
Click Finish to start creation

Provisioning type	Deploy from template
Name	TNEPPDM
Template name	powerprotect
Download size	7.2 GB
Size on disk	600.9 GB
Folder	DFTrekDataCenter
Resource	obys28.all.lab.emc.com
Storage mapping	1
All disks	Datastore: datastore1 (1) Format: Thick provision lazy zipped
Network mapping	1
VM Network	VM Network
IP allocation settings	
IP protocol	IPv4
IP allocation	Static - Manual

CANCEL BACK FINISH

Figure 2. PowerProtect Data Manager OVF deployment

Post-installation steps for PowerProtect Data Manager

Once the PowerProtect Data Manager OVA is successfully deployed and the VM is powered on, you can access the PowerProtect Data Manager UI. Enter the IP address or FQDN configured during the deployment using `https://appliance_hostname`.

The security certificate that encrypts communication between the PowerProtect Data Manager UI and the web browser is self-signed. A self-signed certificate has been signed by the web server that hosts the secure web page being viewed by a web browser. This

certificate is enough to establish an encrypted channel between the web browser and the server. However, it has not been signed by a trusted authority.

The next step is to configure the basic settings for the PowerProtect Data Manager. Google Chrome is the only supported browser.

Welcome screen The **Welcome** screen is displayed when the PowerProtect UI is accessed for the first time after deploying the OVF template. There are two options available on this screen: **New Install** and **Restore Backup**.

- **New Install:** If you are installing the device for the first time, click **New Install**.
- **Restore Backup:** This option is used when PowerProtect Data Manager appliance must be restored from a previous backup. To delay jobs defined by your protection policies until otherwise specified, select the option **After restoring, keep the product in recovery mode so that scheduled workflows are not triggered**. A system alert is displayed in the PowerProtect Data Manager.



Figure 3. PowerProtect Data Manager welcome screen

End-user license agreement (EULA) Scroll through the agreement and accept the license.



Figure 4. PowerProtect Data Manager EULA

License PowerProtect Data Manager has three types of license-trial license (valid for 90 days), Front-End Terabyte (FTEB) protected capacity and socket-based.

The relevant .xml license file can be obtained from the Dell license management website. To obtain the license file, you must have the License Authorization Code (LAC), which was emailed from Dell. If you have not received the LAC, contact your technical support representative.



Figure 5. PowerProtect Data Manager license type selection

Authentication The authentication screen allows the administrator to set an authentication password for appliance management. In addition, the administrator can set the same password or a different one using the toggle key option for the root, administrator, and support accounts for Linux, SSH, and support activities.

From the PowerProtect Data Manager CLI, password expiry can be set to never for users such as administrator, root, and support accounts by running the following command:

```
chage -m 0 -M 99999 -I -1 -E -1 <role-name>
```

- The administrator password is used to log in to the PowerProtect Data Manager management console.

Note: PowerProtect Data Manager application user's password is set to 60 days by default. To extend the expiration time, update the value of the parameter `aaa.server.policies.password.maxAge` in the properties file: `/usr/local/brs/14b/aaa/config/application-policies-custom.properties`

- The service password is used to log in to SSH for support activities.
- The lockbox password is used during restore operations.

Note: You can select **Use of the same password for all** to set the same password for administrator, service, and lockbox accounts, but it is not recommended. Save all the passwords securely for later use.

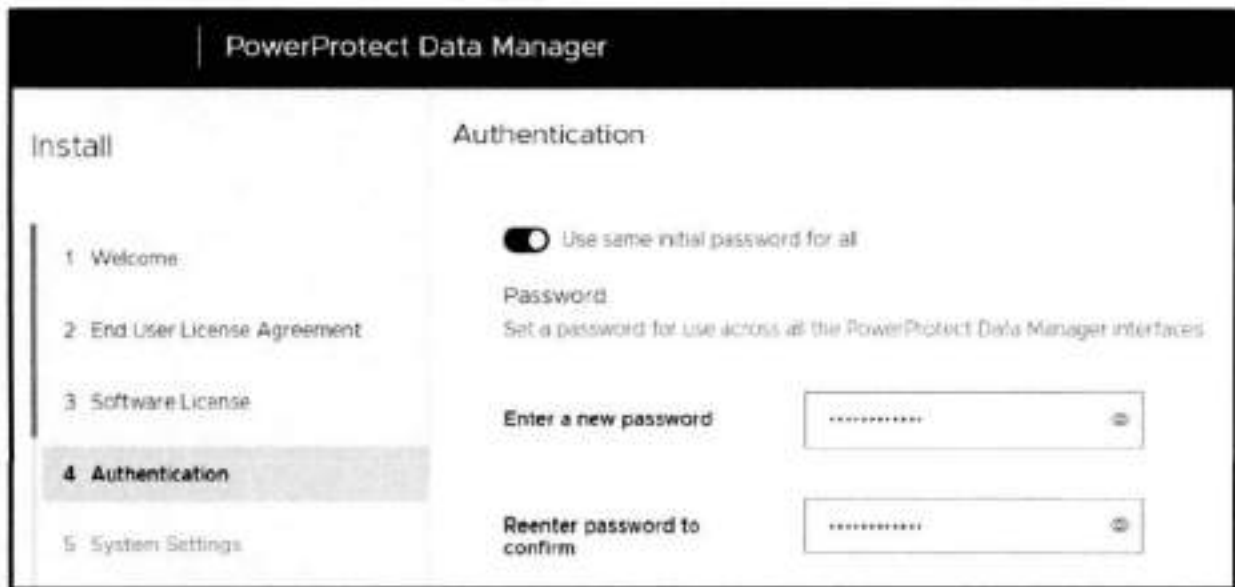


Figure 6. PowerProtect Data Manager authentication

System settings The system settings screen allows the user to add the NTP server. Dell Technologies recommends that users always add an NTP server and sync the PowerProtect Data Manager with it.



Figure 7. PowerProtect Data Manager system settings

Email and SMTP setup Email setup is an optional step. The SMTP server can be configured to receive the PowerProtect Data Manager alerts using a specified email address. To send diagnostic and usage data to Dell for proactive support and to help improve our products and services, switch Auto Support to ON.

Post-installation steps for PowerProtect Data Manager

The screenshot displays the 'Email Setup' step of the installation wizard. On the left, a vertical list of steps is shown, with '6 Email Setup' highlighted. The main area contains the following fields:

- Mail Server:** A text input field containing 'smtp.gmail.com' with a '(Required)' label to its right.
- Admin Email:** A text input field containing 'admin@powerprotect.com' with a '(Required)' label to its right.
- Recipient for Test Email:** A text input field containing 'john@gmail.com' with a '(Required in order to send a test email)' label below it.
- Port:** A text input field containing '25'.
- Authentication:** A section header for the following fields.
- User Name:** A text input field containing 'admin'.
- Password:** A password input field with a masked password '*****' and a visibility toggle icon.

Figure 8. PowerProtect Data Manager email setup (optional)

Login and Getting Started

Once the DONE button on the summary screen is selected, the PowerProtect Data Manager applies all the configured settings, and the login screen is displayed after the setup is complete.

A user can now log in to the PowerProtect Data Manager using the credentials specified in the authentication step. The default login username is `admin`.



Figure 9. PowerProtect Data Manager configuration in progress



Figure 10. PowerProtect Data Manager login screen

After logging in, the **Get Started** screen is visible. The PowerProtect Data Manager returns to the **Get Started** screen until you click **Skip This**.

The **Get Started** screen can be accessed at any time through **System Settings > Getting Started**. The navigation options on the **Get Started** screen are:

- **License:** Links to the License addition page
- **Support:** Links to the SupportAssist configuration page
- **Assets:** Links to the Asset sources page.
- **Storage:** Links to the target Storage Addition or configuration page

Click **Launch** to skip this step and go to the main UI.



Figure 11. PowerProtect Data Manager Get Started screen

Configure 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 Technical Support for diagnostic purposes and Customer Support.

SupportAssist provides the following features and benefits:

- Proactive monitoring and issue prevention
- Facilitates upgrade package downloads
- Automatic support case creation based on event alerting
- Automatic health checks
- Communicates telemetry data
- Real-time troubleshooting
- Customer support

Note: SupportAssist provides automated support capabilities for PowerProtect Data Manager systems. SupportAssist replaces Secure Remote Services in the current release of PowerProtect Data Manager. If you have configured Secure Remote Services previously, the PowerProtect Data Manager system automatically migrates Secure Remote Services to SupportAssist when you upgrade the PowerProtect Data Manager. SupportAssist cannot be configured when you have a trial license.

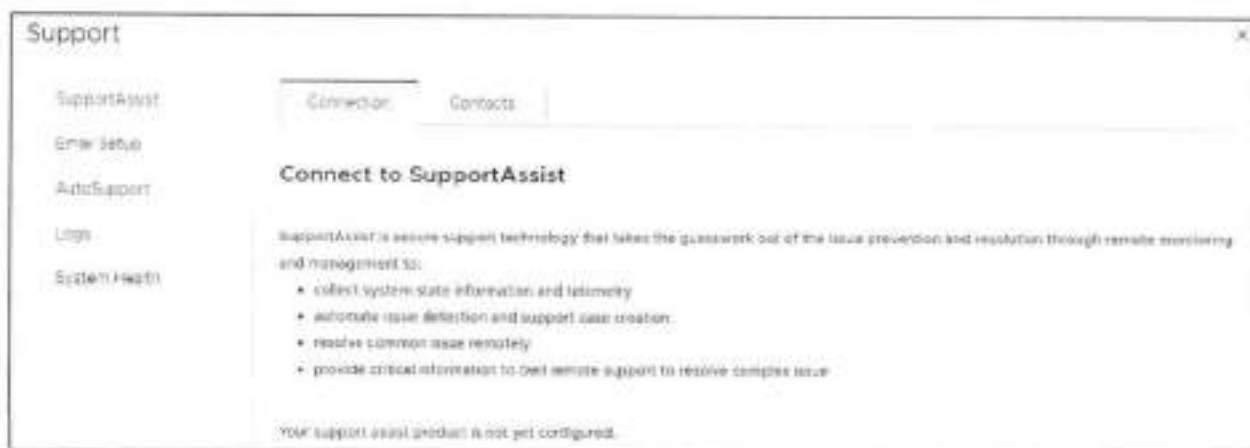


Figure 12. Connecting to SupportAssist

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

For details, see [PowerProtect Data Manager Deployment Guide](#) to complete this setup.

Setting up disaster recovery of PowerProtect Data Manager

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 system backups.

- Each backup is considered a “full” backup although it is created in an incremental manner. The persistent data that is saved in a backup includes the Lockbox and Elasticsearch databases.
- The backup operation creates a Point-in-Time snapshot of the database while the system is in a quiesced state. While the system is quiesced, user functionality is limited. After the snapshot is completed, and while PowerProtect Data Manager copies the snapshots to the DD storage unit, full user functionality is restored. If the system fails to quiesce, PowerProtect Data Manager still takes a backup, which is marked as crash consistent instead of application consistent.
- To store system backups, you must configure and assign a private DD storage unit for the PowerProtect Data Manager system. The system protection service enables you to manage the frequency and start time of an automated system backup, perform manual backups, and define the length of time that the system backups are available for recovery.

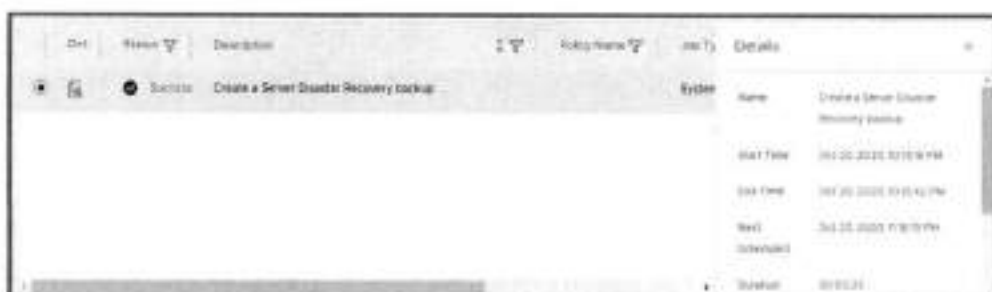


- File Search indexes are backed up for DR recovery along with other component DR backups.
- Make the following selections, and then click **Save**.
 - Select **Enable Backup**.
 - At **Protocol**, select **NFS** or **DDBoost**. (DDBoost is recommended.)
 - At **PowerProtect DD System**, select the system from the drop-down menu if it has already been added; otherwise, add a PowerProtect DD appliance.

Post-installation steps for PowerProtect Data Manager



The initial backup runs, and then backups are automatically triggered every hour.



Add asset sources

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 storage system where backup copies of the assets are stored.

Asset sources can be a vCenter Server, Kubernetes cluster, application host, or SMIS server. Assets can be Virtual Machines, Exchange databases, SQL databases, Oracle databases, SAP HANA databases, File systems, Kubernetes namespaces, or Storage Groups.

Add a vCenter Server

An asset source, such as a vCenter Server, must be enabled in PowerProtect Data Manager before you can add and register the asset source for the protection of assets.

Follow these steps to add a vCenter Server as an asset source in the PowerProtect Data Manager UI:

1. Select **Infrastructure > Asset Sources > Virtual Machine > Enable Source**.



2. Enter the vCenter details and click **Save**.



The initial vCenter Server discovery identifies all VMware ESXi clusters, hosts, and virtual machines within the vCenter Server. Subsequent discoveries are performed automatically according to a fixed interval to identify any additional or changed VMware entities since the last discovery operation. You can also manually initiate a discovery of VMware entities

at any time from the **vCenter** tab of the **Asset Sources** window by selecting a vCenter Server and clicking **Discover**.

After the vCenter virtual machine assets are discovered, you can add a VM Direct appliance to facilitate data movement and then create virtual machine protection policies to back up these assets.

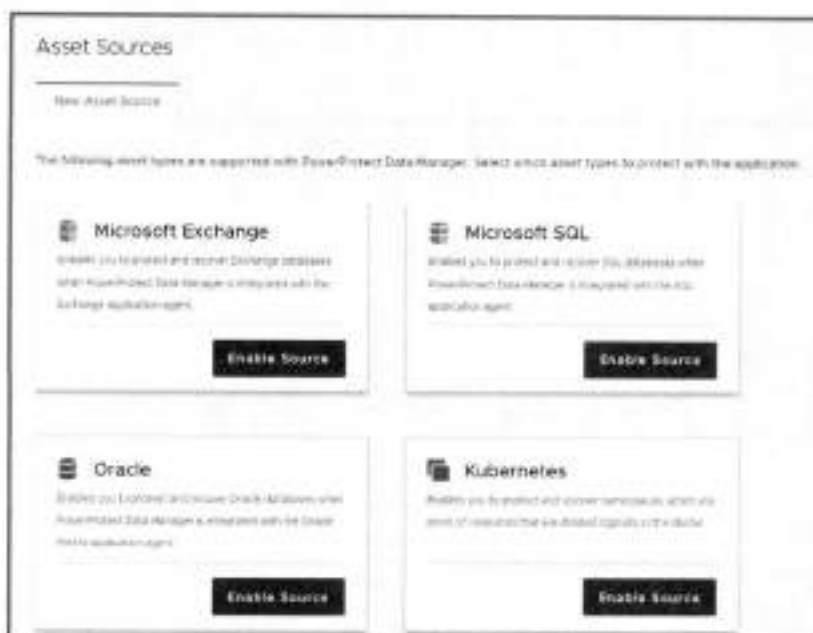
Add other asset sources

In addition to vCenter Server asset sources, PowerProtect Data Manager provides the option to enable the following asset sources to protect application assets.

- Kubernetes Cluster
- File System Agent
- Microsoft Exchange Agent
- Microsoft SQL Agent
- Oracle RMAN Agent
- SAP HANA Agent
- Storage Direct Agent for Storage Data Management

Note: This white paper does not provide instructions for each application agent. For links to the individual application agent user guides, see [User permission requirements for supported platforms](#).

Select **Infrastructure** > **Asset Sources** > **Enable Source**.



Configuring PowerProtect Search Engine

When you install PowerProtect Data Manager version 19.3 or later, the PowerProtect Search Engine is installed by default. The following bullet points explain its usage:

- 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 search cluster, and then enable the indexing feature.
- PowerProtect Search 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. When Search is enabled, you must allow the Search Engine virtual machine on the DD series appliance that contains the Server Backup MTree: Add the search node IP address or hostname to the client list for the NFS export.

Add and configure the Search Engine as follows:

1. To add the Search Engine, select **Infrastructure > Search Engine** and click **Add Node**.



2. In the **Add Search Engine Node** dialog box, enter the required network parameters.

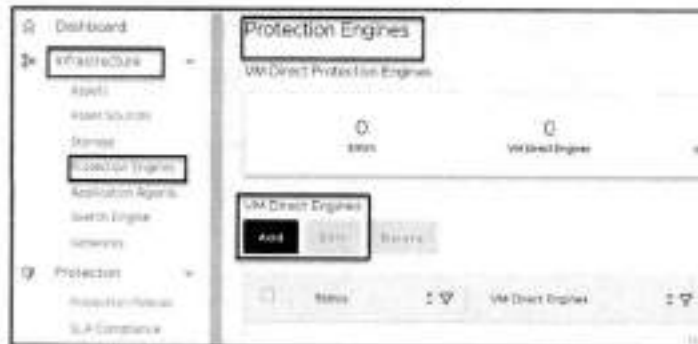
3. Select **Add** and wait a few minutes for the VM to show the **Node State** as **Ready**.

Post-installation steps for PowerProtect Data Manager



Adding external VM Direct Engine In the **Protection Engines** window, deploy an external VM Direct Engine, also referred to as a VM proxy, to facilitate data movement for virtual machine protection policies:

1. Select **Infrastructure > Protection Engine > Add**.



2. Enter the network details and click **Next**.



3. Select the VLAN network configured for the VM Direct Engine.

For more details about adding multiple VLANs, see

Multiple VLAN configurations.



The proxy is automatically deployed on the vCenter. Once its status shows that it is ready, you can configure VM backups.



Adding PowerProtect DD series appliance

The PowerProtect Data Manager UI enables users with administrator credentials to add the following storage types:

- PowerProtect DD Management Center
- External PowerProtect DD series appliance

For each PowerProtect DD series appliance, the PowerProtect DD Management Center that manages the DD system is indicated in the **Managed By** column in the table.

If a DD series appliance is added directly to the PowerProtect Data Manager, the name that was provided for the DD series when it was added to the PowerProtect Data Manager system is displayed in the **Managed By** column.

Add the PowerProtect DD series appliance:

1. Select **Infrastructure > Storage** and click **Add**.

Post-installation steps for PowerProtect Data Manager



2. Enter the DD series appliance details and click **Save**.



3. Add the host credentials and click **Save**.



4. At **Verify Certificate**, click **Accept**.5. Under **Add Storage**, confirm that the certificate status is displayed as **Verified**, and then click **Save**.**Verification**

Discovery is completed within a few minutes. Verify that DD series appliance has been added to the PowerProtect Data Manager successfully by selecting **Infrastructure > Storage > Protection Storage** tab.



You are now ready to start configuring the backups. For more information about configuring VM, file system, and application backups, along with other management activities, see the [PowerProtect Data Manager Administration and User Guide](#).

Multiple VLAN configurations

Multiple VLAN configurations: Introduction and prerequisites

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.
- Virtual networks can also be configured to separate management traffic from backup traffic. This configuration can also separate traffic that originates from different networks.

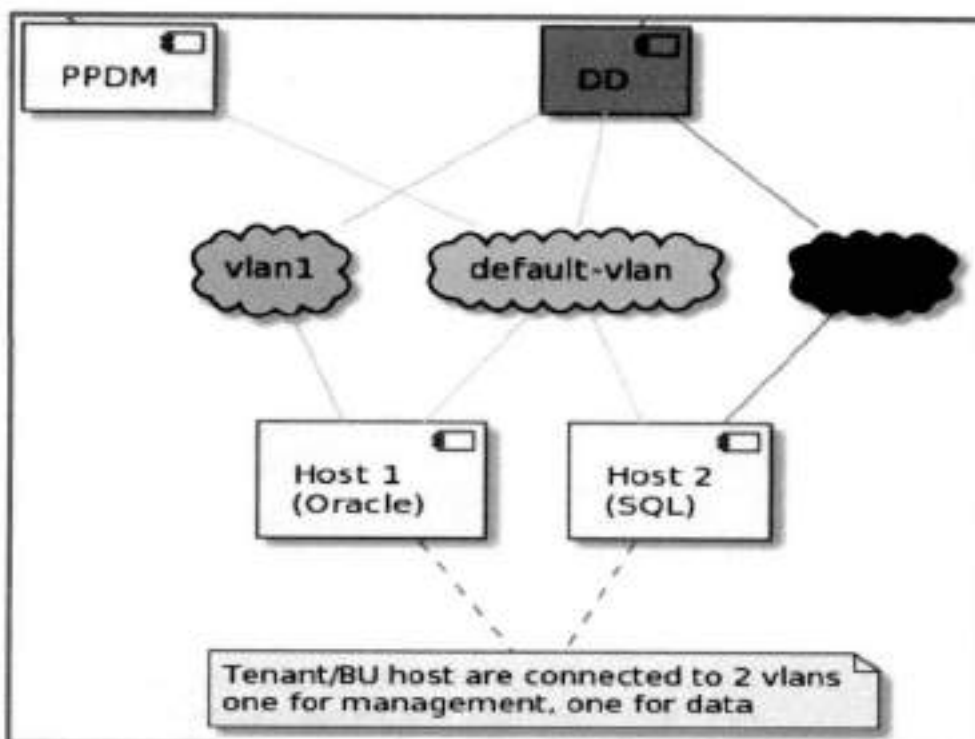


Figure 13. Example of 2-VLAN environment

Before you configure a virtual network, complete the following actions:

1. Register the vCenter server on which the PowerProtect Data Manager is deployed. You can verify the registration on the **vCenter** tab of the **Asset Sources** page.
2. Configure the network switch port for trunk mode. This setting allows the port to carry traffic for multiple VLANs.
3. Enable Virtual Guest Tagging (VGT) mode on the VMware ESXi virtual network switch port for the PowerProtect Data Manager. Configure the virtual switch port for VLAN ID 4095.
4. 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 relating to this activity.

5. Add the DD series appliance as protection storage for the PowerProtect Data Manager.

Note: PowerProtect Data Manager does not verify the network switch configurations. If the physical or virtual network switch is incorrectly configured, the virtual network configuration fails.

Configuring VLAN

PowerProtect Data Manager names each virtual network in two places, the interface to the DD series appliance and the interface to the protected assets. These names are not required to match. However, Dell Technologies strongly recommends that you use the same network name in both locations for each virtual network. Record each network name for later use.

- 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 Engines that you deploy on this network. 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.
- 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.

Configuration follows a multistep workflow as follows:

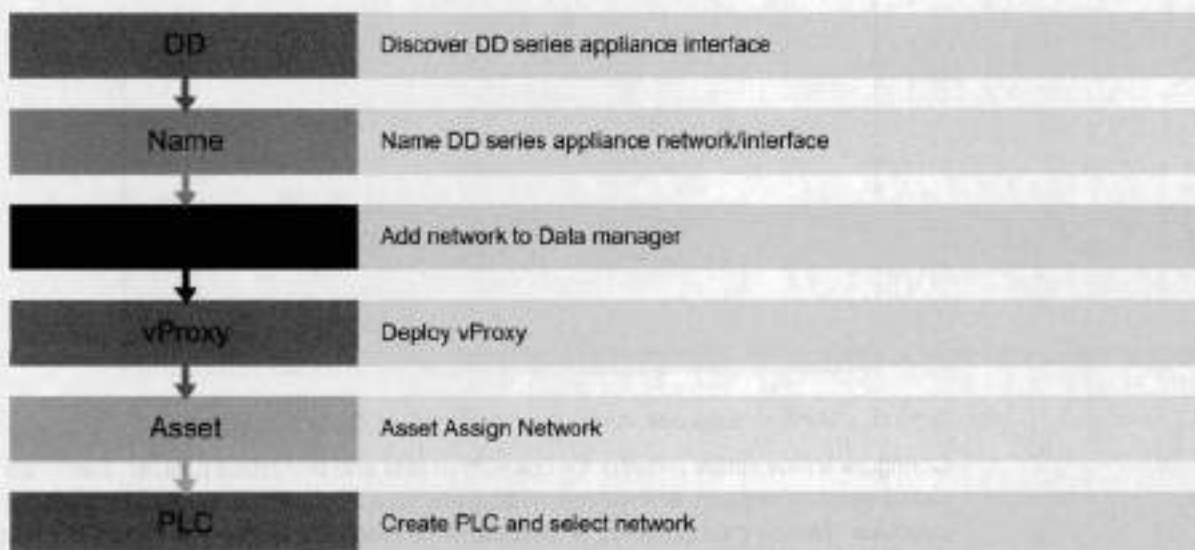


Figure 14. VLAN configuration steps

Discover and name PowerProtect DD series network or interface

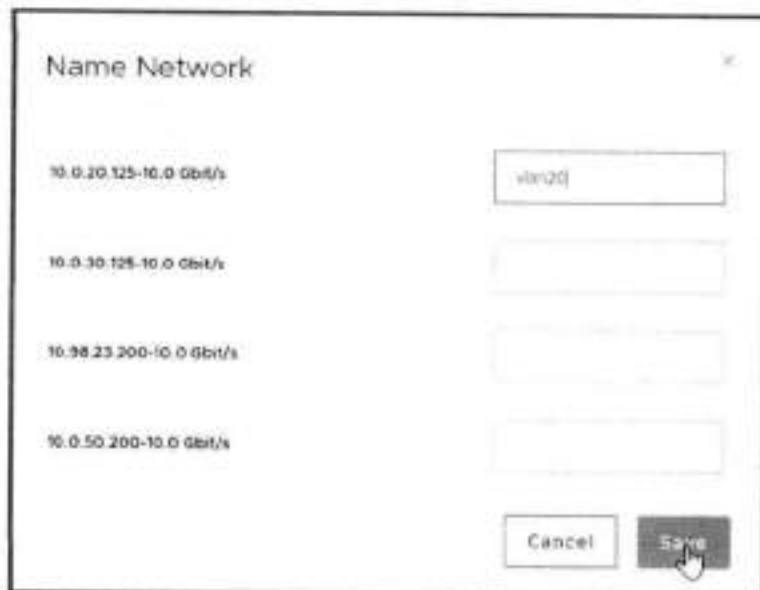
After adding the DD series appliance as protection storage, name the virtual network between the PowerProtect Data Manager and the DD series appliance.

To rename a virtual network (edit the network name):

1. Select **Infrastructure > Storage > Protection Storage** and select the **Name Network** tab.



2. Select the network, add the VLAN name, and click **Save**.



Add the virtual network to the PowerProtect Data Manager

Configure a new virtual network for use with assets and protection policies. Each new virtual network requires at least one IP address for a PowerProtect Data Manager network interface. Review the number of IP addresses needed before you supply the required static IP addresses.

Select **Infrastructure > Networks** and click **Add Network**.



Assign the preferred virtual network to a protection policy or asset

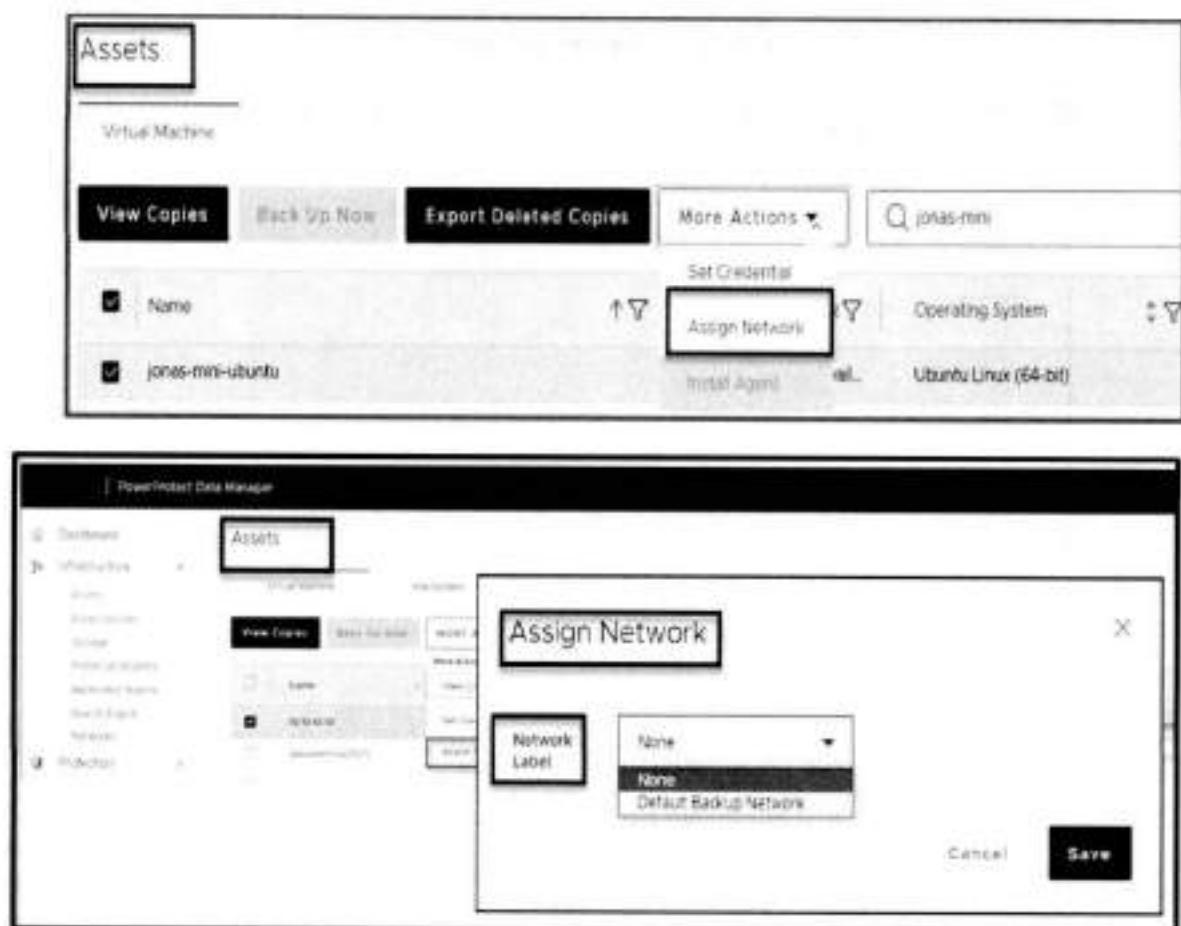
Assignments identify which assets should use each virtual network. You can associate an asset with a virtual network by using one of two methods:

- Using protection policy:** PowerProtect Data Manager can be configured to choose a preferred virtual network for all assets on a protection policy. The network interface has a drop-down menu, and you can select the preferred network for Primary backup and Replicate.



- Using asset:** Virtual networks can be assigned to individual assets. This method is optional and overrides any virtual network assignment from a protection policy. Assets that are not individually assigned a virtual network automatically use the preferred virtual network.

Multiple VLAN configurations



Supported scenarios

PowerProtect Data Manager 19.11 supports virtual networks for the following use cases:

- Virtual machine backups
- Database backups
- Exchange backups
- File system backups
- Replication
- Disaster recovery
- Cloud DR
- Storage data management

Notes and limitations of multiple VLANs

Consider the following information when you are using multiple VLANs:

- PowerProtect Data Manager supports only VGT mode (multi VLAN) unlike vProxy, which supports both VST (single VLAN) and VGT modes. Therefore, the UI does not show a selection for a port group for PowerProtect Data Manager.
- No restrictions apply to the operation flow sequence. Any parameter can be edited to modify or add values for the defined parameters.

- The DD series appliance network name can be different, but PowerProtect Data Manager and assets must reach the DD series appliance network. Using a proper naming convention helps.
- Customers are responsible for ensuring that networks are reachable and configured correctly.
- An old vProxy cannot be edited and attached to a VST or VGT port group. It must be reconfigured.
- We recommend that you have PowerProtect Data Manager and the DD series appliance in the same VLAN; otherwise, proper gateway/routing must be established.
- Search can only be done in the default VLAN.
- PowerProtect Data Manager asset restriction (application and file system): If the assets of the same host/client are attached to different networks, those assets must be on different policy.

Scalability limits for PowerProtect Data Manager

The following limits have been tested successfully with PowerProtect Data Manager for the vCenter Server, the VM Direct Engine, and DD systems.

Table 7. Tested limits for PowerProtect Data Manager components

Component (per PowerProtect Data Manager)	Tested limits
vCenter Servers supported	12
External VM Direct Engines supported	40
DD series appliances supported	10
Virtual machines	10,000
Maximum search engine node	5

Notes:

These numbers are not maximum (hard) limits but should be considered as best practice when scaling your environment.

The vCenter server limit is subject to the VM Direct Engines overall limit of 40 and per vCenter limit of 25. For example, using the maximum tested number of vCenter servers (12), you could add an average of three VM Direct Engines per vCenter.

The number of external VM Direct Engines was tested across 10 vCenter servers (for example, 4 VM Direct Engines per vCenter).

Conclusion

This document provides a detailed overview of the PowerProtect Data Manager deployment requirements, process, and best practices to successfully deploy a new PowerProtect Data Manager.

Technical support and resources

Technical support

For the most up-to-date software compatibility information for PowerProtect Data Manager, see the E-Lab Navigator at <https://elabnavigator.emc.com/eln/modernHomeDataProtection>.

For technical support, see the [Dell Support](#) website.

Dell Technologies documentation

The following links provide other information related to this document. Access to documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

- [PowerProtect Data Manager Info Hub](#)
- [PowerProtect Data Manager Administration and User Guide](#)
- [PowerProtect Data Manager Deployment Guide](#)
- [PowerProtect Data Manager Security Configuration Guide](#)
- [PowerProtect and Data Domain core documents](#)

Dell PowerProtect Data Manager: Oracle RMAN Agent Backup and Recovery

Abstract

This white paper focuses on protecting an Oracle database using Dell PowerProtect Data Manager, the next-generation data protection platform.

October 2022

Revisions

Date	Description
July 2019	Initial release
February 2021	Revised
May 2021	Revised
July 2022	Revised for PowerProtect Data Manager version 19.11 release
October 2022	Revised for PowerProtect Data Manager version 19.12 release

Acknowledgments

Author: Vinod Kumar Kumaresan and Chetan Padhy

The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

This document may contain certain words that are not consistent with Dell's current language guidelines. Dell plans to update the document over subsequent future releases to revise these words accordingly.

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

Copyright © 2019–2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [10/30/2022] [White Paper] [H18626.4]

Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents.....	3
Executive summary.....	5
Audience.....	5
Introduction.....	6
1.1 Solution components.....	6
1.2 PowerProtect DD series appliances.....	6
1.3 PowerProtect Data Manager.....	7
1.4 Oracle RMAN agent.....	8
1.5 PowerProtect Data Manager agent.....	8
1.6 Oracle Recovery Manager server.....	8
Installation of the Oracle RMAN agent.....	9
1.7 Deployment requirements.....	9
1.8 Install and configure Oracle RMAN agent.....	9
1.8.1 Roadmap to protect Oracle database with PowerProtect Data Manager.....	9
1.8.2 Setting Oracle RAC Preferred Node Using Data Manager.....	12
2.2.2.1 Asset Details.....	12
1.9 Authentication requirements.....	13
1.9.1 Setting Oracle assets credentials in Data Manager.....	13
1.10 Verification of database connectivity.....	15
1.10.1 System verification.....	15
1.10.2 Asset verification.....	15
1.10.3 RMAN verification.....	15
PowerProtect Data Manager protection policy.....	15
1.11 Centralized protection policy.....	16
1.12 Self-Service protection policy.....	17
1.13 Storage unit consideration.....	18
1.14 Top-level directory changes.....	18
Oracle RMAN agent backup workflow.....	20
1.15 Centralized protection backup.....	20
1.15.1 Backup levels for centralized protection.....	21
1.15.2 Enable multistream backup.....	22
1.15.3 Archive log backup.....	23

1.15.4	Monitoring jobs and task for centralized protection policy	24
1.16	Self-service protection backup	25
Oracle RMAN	recovery	28
1.17	Centralized restore and recovery of Oracle backups	28
1.17.1	Centralized Oracle restore and recovery of a full online database	29
1.17.2	Centralized Oracle restore of archive logs	31
1.17.3	Centralized disaster recovery of an Oracle database	34
1.18	Self-service restores of Oracle databases	38
Oracle Data Guard	support	40
1.19	Data Guard configuration	40
1.20	Data Guard: standalone mode support	41
1.21	Data Guard: the recovery catalog option	41
1.22	Data Guard: the disaster recovery option	42
1.23	Data Guard: self-service protection	42
Replication and DD Cloud Tier	43
A	Technical support and resources	46
A.1	Related resources	46

Executive summary

Dell 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 enables the transformation from traditional centralized protection to an IT-as-a-service model based on a self-service design. This design ensures that you can enforce compliance and other business rules, even when backup responsibilities are decentralized to individual database administrators and application administrators. Data Manager key features include:

- 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

The modern, services-based architecture provides ease of deployment, scaling, and upgrading of Data Manager and integrates multiple data protection products within the Dell Data Protection portfolio to enable data protection as a service. Data Manager enables the protection of traditional workloads including Oracle, Exchange, SQL, SAP HANA, and file systems as well as Kubernetes containers and virtual environments.

Audience

This white paper is intended for customers, partners, and employees who want to better understand, evaluate, and explore Data Manager Integration with Oracle RMAN agent. Familiarity with PowerProtect Data Manager and DD series appliances is required.

Introduction

Data Manager integrates with the Oracle RMAN agent to check and monitor backup compliance against protection policies. The Oracle RMAN agent enables an application administrator to protect and recover the Oracle data on the application host.

1.1 Solution components

This section discusses PowerProtect Data Manager and Oracle RMAN agent backup components.




Figure 1 PowerProtect Data Manager Oracle RMAN agent solution components overview

1.2 PowerProtect DD series appliances

Dell PowerProtect DD series appliances and older Data Domain systems are disk-based appliances that run DDOS to provide inline deduplication for data protection and disaster recovery (DR) in the enterprise environment. DD series appliances vary in storage capacity and data throughput. Systems are typically configured with expansion enclosures that add storage space.

DDOS features include:

- **Data integrity:** The DDOS Data Invulnerability Architecture protects against data loss from hardware and software failures.
- **Data Deduplication:** The file system deduplicates data by identifying redundant data during each backup and storing unique data once.
- **Restore operations:** File restore operations create little or no contention with backup or other restore operations.
- **DD Replicator:** DD Replicator sets up and manages the replication of backup data between two protection systems.
- **Multipath and load balancing:** In a Fibre Channel multipath configuration, multiple paths are established between a protection system and a backup server or backup destination array. When multiple paths are present, the system automatically balances the backup load between the available paths.
- **High availability:** The High Availability (HA) feature lets you configure two protection systems as an Active-Standby pair, providing redundancy in the event of a system failure. HA keeps the active and standby systems that are synchronized, so that if the active node were to fail due to hardware or software issues, the standby node can take over services and continue where the failing node left off.

- **Random I/O handling:** The random I/O optimizations in DDOS provide improved performance for applications and use cases that generate larger amounts of random read and write operations than sequential read and write operations.
- **System Administrator access:** System administrators can access the system for configuration and management using a command-line interface (CLI) or a user interface (UI).
- **Licensed features:** Feature licenses allow you to purchase only those features you intend to use. Some examples of features that require licenses are DD Boost, and capacity on demand (storage capacity increases).
- **Storage environment integration:** DDOS systems integrate easily into existing data centers.

1.3 PowerProtect Data Manager

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

Data Manager enables the transformation from traditional centralized protection to an IT-as-a-service model based on a self-service design. This design ensures that you can enforce compliance and other business rules, even when backup responsibilities are decentralized to individual database administrators and application administrators.



Figure 2 PowerProtect Data Manager highlights

Data Manager key features include:

- 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 upgrading

1.4 Oracle RMAN agent

The Oracle application agent allows an application administrator to protect and recover the Oracle RMAN application data on the application host. Oracle RMAN agent can be used alone by DBA to backup using RMAN scripts and transfer the data to DD series appliance. Data Manager agent integrates with the RMAN agent to check and monitor backup compliance against protection policies. Oracle RMAN agent has two main components:

- **ddbmcon**: The ddbmcon program is installed with the Oracle RMAN agent software and enables the Data Manager monitoring, management, and analysis of Oracle RMAN agent backups. During the Oracle RMAN agent installation, the ddbmcon program is installed in the \$RMAN_HOME/bin directory. You cannot run the ddbmcon program manually. The program is only run by the PowerProtect Data Manager agent.
- **ddutil**: The ddutil program is also installed with the Oracle RMAN agent software. It enables DD series appliance command-line interface, which is used to create the lockbox, verify connectivity with DD series appliance, display or delete a backup, and perform various other manual operations that can be useful during self-service backup. For more details about command-line options, see the document [PowerProtect Data Manager for Oracle RMAN Agent User Guide](#).

1.5 PowerProtect Data Manager agent

PowerProtect Data Manager agent is installed with Oracle RMAN agent. With the help of PowerProtect Data Manager agent administrators can monitor, manage, or analyze the Oracle RMAN agent backups on AIX or Linux. PowerProtect Data Manager agent can create and manage replication copies based on the protection policies. PowerProtect Data Manager agent performs these operations whether the backup is created by the DBA or by the Data Manager centralized backup scheduler. Data Manager with Oracle integration supports the following features and functionalities:

- Data Manager Centralized protection policy.
- Data Manager Self-Service protection policy.
- Oracle databases can be discovered by Data Manager automatically.
- Single Protection Policy (PLC) to manage the entire life cycle of data protection for Oracle database
- Centralized Oracle restore and recovery from the PowerProtect Data Manager UI
- Support of Oracle RMAN agent with Oracle 21c by the current version of PowerProtect Data Manager and the two previous versions. See [PowerProtect Data Manager Compatibility Matrix](#) for more details.

1.6 Oracle Recovery Manager server

Recovery Manager (RMAN) is an Oracle utility that can backup, restore, and recover database files. The product is a feature of the Oracle database server and does not require separate installation.

Recovery Manager is a client/server application that uses database server sessions to perform backup and recovery. It stores metadata about its operations in the control file of the target database and, optionally, in a recovery catalog schema in an Oracle database. You can invoke RMAN as a command-line executable from the operating system prompt or use some RMAN features through the Enterprise Manager UI.

Installation of the Oracle RMAN agent

Data Manager can manage and monitor data protection and replication for Oracle assets through integration with the Oracle RMAN agent. For configuring Data Manager with Oracle RMAN agent, the following two software components must be installed on the Oracle database host:

- Oracle RMAN Agent
- PowerProtect Data Manager agent

In an Oracle Real Application Clusters (RAC) environment, the Oracle RMAN agent and the PowerProtect Data Manager agent must be installed on each node in the Oracle RAC environment.

1.7 Deployment requirements

Ensure that you meet the prerequisites before you add an Oracle asset.

Verify that the environment meets the following requirements:

- Ensure that all clocks on both the Oracle host and Data Manager are time-synced to the local NTP server to ensure discovery of the backups
- Ensure that the Oracle host and the Data Manager network can see and resolve each other
- Ensure that port 7000 and 8443 is open on the Oracle host. Port 111,2049 and 2052 are open between DD series appliance and Oracle hosts
- Ensure that DD series appliances are added as the protection storage

1.8 Install and configure Oracle RMAN agent

1.8.1 Roadmap to protect Oracle database with PowerProtect Data Manager

Data Manager can manage and monitor data protection and replication for Oracle assets through integration with the Oracle RMAN agent.



Figure 3 Oracle RMAN agent install and configure workflow

Oracle RMAN agent can be download from **Dashboard > System Settings> Downloads**.

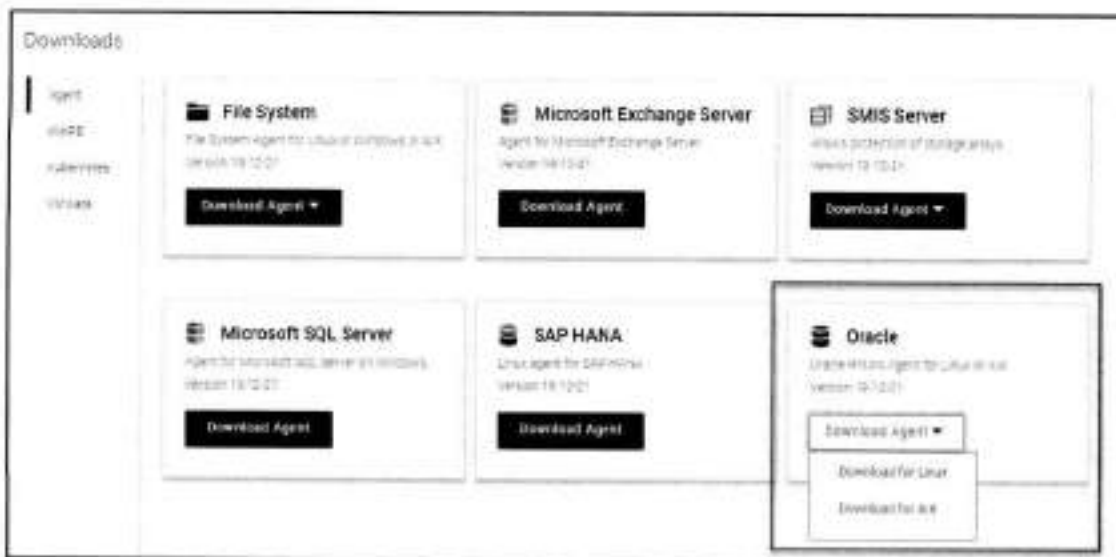


Figure 4 Oracle RMAN agent download

Use this option to add or approve the Oracle RMAN agent from **Infrastructure > Application Agents**.



After you register an application host with PowerProtect Data Manager, you can use the **Asset Sources** window to discover an application host and modify the application host credentials. You must add credentials to the Oracle database so that PowerProtect Data Manager can access the database to create backups.



Note: Starting with release 19.9, a new Oracle database discovery method is supported, which uses the `prons` process without a dependence on `/etc/oratab` entries. The `/etc/oratab` file entries have the highest precedence for the discovery of Oracle database resources on the system, which enables the PowerProtect Data Manager operations.

Once the asset discovery is complete, the Oracle database assets are discovered in the **Infrastructure > Assets** section.



Figure 5 Oracle assets section

See the [PowerProtect Data Manager for Oracle RMAN Agent User Guide](#), for more details about how to install and configure Oracle RMAN agent with PowerProtect Data Manager.

1.8.2 Setting Oracle RAC Preferred Node Using Data Manager

Starting with Data Manager 19.12, you can set the Oracle RAC preferred node in the Data Manager user interface. This applies to a single/multi-asset based Oracle RAC environment.



Note: The Preferred Node value from the config file gets migrated when the Data Manager server is upgraded from an older version. Adding/removing RAC nodes will automatically adjust the configuration.

The following are Oracle RAC preferred node UI features:

- The preferred node can be set with different AUTHs
- Any change in Preferred Node will take effect at the next triggered backup
- Preferred Node from UI takes precedence.
- Decommissioning the preferred node will force Data Manager to use any other available node
- Shutting down the database or stopping the RMAN agent service on the preferred node will result in backup failure
- Works with both Linux and AIX platforms
- Compatible with older RMAN agent

2.2.2.1 Asset Details

After setting the Oracle preferred node in the Data Manager user interface, you can view which host is set as the Oracle RAC Preferred Node from Oracle Asset details.



1.9 Authentication requirements

The Oracle RMAN agent program `ddbmcon` handles all communication between the Oracle RMAN agent and Data Manager. When the `ddbmcon` program performs discovery, backup, or deletion operations, it connects to the Oracle database. The following authentication methods are supported:

1. Database authentication
2. Oracle wallet authentication
3. Operating system authentication

The `ddbmcon` program tries all these authentication methods for each Oracle database instance. The program reports a connection error if it cannot connect to the database instance by using any of these methods. If one of these methods succeeds, the `ddbmcon` program ignores the other authentication methods and goes to retrieve the information as used by the Data Manager.

Ensure that you enable one of these three authentication methods for the `ddbmcon` program. For maximum ease of use, it is recommended that you enable the operating system authentication method. Both the database and Oracle wallet authentication methods require additional configuration steps on the Oracle host and parameter settings in the configuration file `rman_agent.cfg`. It is installed in the `$RMAN_AGENT_HOME/config` directory.

After you have set the authentication method, you can select the same in Data Manager UI during discovery and Policy Lifecycle (PLC) creation. For detailed instructions, see the [PowerProtect Data Manager for Oracle RMAN agent User Guide](#).

1.9.1 Setting Oracle assets credentials in Data Manager

Starting with Data Manager 19.7, you can optionally add and remove the credentials for one or more Oracle database assets simultaneously in the Data Manager UI. Following are few important points regarding adding credentials:

- You can only add the asset-level credentials when the Oracle host agent version is 19.7 or later.
- You can add Oracle assets with different Oracle operating system users or groups from the same asset source into a single protection policy.
- You can add multiple Oracle assets from multiple asset sources into a single protection policy.
- The Oracle assets can be associated with multiple credential types, where the supported database credential types are Oracle, Database User, and Wallet and the supported RMAN catalog credential types are Database User and Wallet.

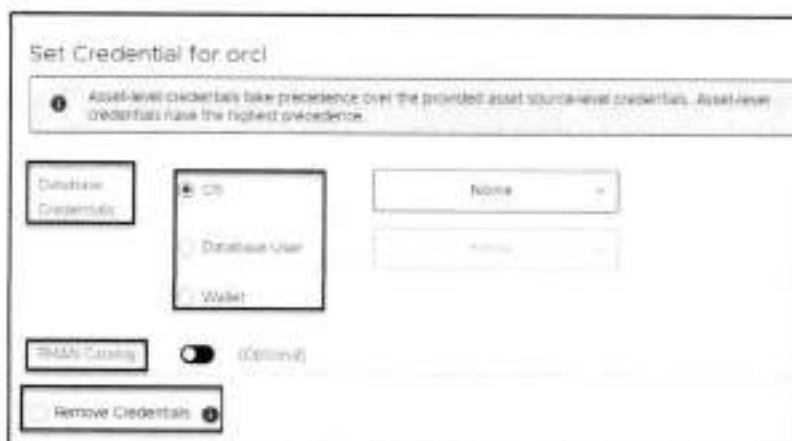
Use the following procedure to add or remove the credentials for the Oracle database assets.

1. In the Data Manager UI, select **Infrastructure > Assets**, and click the **Oracle** tab.
2. Select one or more assets by clicking the checkbox next to each required asset name.
3. Select **More Actions > Set Credential**.



4. In the **Set Credential** dialog box, add or remove the credentials for the selected Oracle assets:

To add the credentials for the assets, specify the required **operating system**, **Database User**, or **Wallet** settings for **Database Credentials**. When the asset is associated with an RMAN catalog, you can also specify the RMAN catalog credentials through the **Database User** or **Wallet** settings for **RMAN Catalog**.



Note: You can specify both the database credentials and RMAN catalog credentials in the Set Credential dialog box.

To remove the credentials for the assets, select **Remove Credentials**.

5. Click **Save** in the **Set Credential** dialog box.

Results

When you save the newly added credentials in the dialog box, Data Manager triggers an autoconfiguration job for the credential update in the respective clients.

After you add the credentials by using this procedure, the asset-level credentials are used for the selected assets during Oracle centralized backups, overriding the policy-level credentials.

Note: Credentials that you set at the asset level and asset source level supersede the credentials that you set at the protection policy level. Credentials at the asset level have the highest precedence.

1.10 Verification of database connectivity

You can run the `ddutil` command as the root user with the appropriate `-v` option to verify the connectivity from the `ddbmcon` program to the Oracle database. For detailed instruction check [PowerProtect Data Manager for Oracle RMAN agent User Guide](#). The following subtopics describe the three supported levels of verification with the `ddutil -v` command:

1.10.1 System verification

To perform the system verification, run the `ddutil -v system` command as the root user.

The `ddutil -v system` command verifies the connectivity to the Oracle instances.

1.10.2 Asset verification

To perform the asset verification, run the `ddutil -v asset` command as the root user.

The command verifies the ability to read the Oracle database objects and provides similar output to the system verification command.

1.10.3 RMAN verification

To perform the RMAN verification, run the `ddutil -v rman` command as the root user. This verification is required only if you use an RMAN catalog.

The `ddutil -v rman` command tests whether the `ddbmcon` program can connect to the target database and catalog database through an RMAN script, as required to perform an active deletion of Oracle backups.

Note: Database authentication or Oracle wallet authentication can be used to connect to an RMAN catalog. Operating system authentication cannot be used with the RMAN catalog. For more details, check [PowerProtect Data Manager Oracle RMAN agent User Guide](#).

PowerProtect Data Manager protection policy

With Data Manager, the Oracle database protection task has been transferred from a backup administrator to the Oracle database administrator (or Oracle database owner). Data Manager creates Oracle database backups and manages remote replication copies based on the Protection Policy (PLC). Data Manager performs the backup and replication operations based on the protection policy and governed by the SLA. Oracle databases can be backed up through:

- Automatic backup by the Centralized Protection policy.
- Manual backup by the Oracle database administrator and governed by the Self-Service Protection policy.



Figure 6 Types of protection policies

Note:

- When you create protection policies for RAC databases, ensure that all nodes in the RAC environment are powered on and registered at the time of the protection policy creation. Otherwise, the protection might fail.
- RAC Node that will do the backup must be set by the DBA after the installation is complete using `IS_RAC_BACKUP_NODE =NODENAME` in `xman_agent.cfg` located in `$RMAN_AGENT_HOME/config`.
- For Oracle Instance Group assets, ensure that the maximum length of the hostname plus storage unit is 59. There are no special character limitations.
- Before you perform a backup on a weekly or monthly schedule from the protection policy, ensure that the Data Manager time zone is set to the local time zone. If the Data Manager time zone is not set to the local time zone, the weekly or monthly backup still runs but is triggered based on the Data Manager time zone.
- If applicable, complete all the virtual network configuration tasks before you assign any virtual networks to the protection policy. The *PowerProtect Data Manager Administration and User Guide* provides more information.

1.11 Centralized protection policy

When Data Manager admin creates a protection policy for Oracle databases, the centralized protection option enables the Data Manager to centrally manage the entire life cycle of data protection operations for the Oracle databases.

The data protection attributes are specified when the centralized protection policy is created: Type, purpose, credentials, assets, schedule, replication schedule, options, and SLA. After the protection policy creation is complete, the lockbox is automatically created for source and replication DD series appliance.

Attributes	Attribute option
Type of application	Oracle database
Purpose of the Protection Policy	Centralized Protection
Application Login Credentials	<Specify or select application login credentials>
Application Assets	<Select the desired Oracle databases>
Schedule	Backup Level

	Retention Period Backup start and end time <Specify or select the desired SLA> Replication schedule
Options	<Advanced options>
Summary	Check the option and Save

1.12 Self-Service protection policy

When Data Manager admin creates a protection policy for Oracle databases, the self-service protection option enables the data owner to perform the manual backup operation from the command-line interface. The Data Manager prepares the environment to accommodate the manual backup operations. A few examples of these operations are creating a DD user with a password, creating a DD storage unit, enforcing the backup data retention. After the protection policy creation is completed, the lockbox is automatically created for source and replication DD series appliance.

The data protection attributes are specified when the self-service protection policy is created: Type, purpose, assets, retention, replication schedule, and SLA. Note that only the retention period and replication schedule can be specified in the schedule attribute in the self-service protection policy.

Attributes	Attribute option
Type of application	Oracle database
Purpose of the Protection Policy	Self-Service Protection
Application Assets	<Select the desired Oracle databases>
Schedule	Retention Period <Specify or select the desired SLA> Replication schedule
Summary	Check the option and Save

Note: Data Manager can create and manage replication copies based on the protection policies. Data Manager performs these operations whether the backup is created by self-service policy or by the centralized backup policy.

Because Data Manager controls the replication, when the Oracle RMAN agent is deployed with Data Manager, the following self-service replication operations are disabled:

- Creation of multiple backup copies with the `RMAN BACKUP COPIES` command.
- MTree replication to create backup copies on a secondary DD series appliance.
- You can restore from replicated copies of backups that were performed with a previous version of Oracle RMAN agent.

- When you perform a self-service backup managed by Data Manager, the Data Manager protection policy settings for the given database override the target protection storage settings that are specified in the RMAN backup script. These settings include the Data Domain server hostname and storage unit name.

1.13 Storage unit consideration

When you create a protection policy, the Data Manager software can either create or reuse a storage unit on the specified DD system backup host, subject to limitations. All subsequent backups of assets in that protection policy go to this storage unit.

The storage unit set using Data Manager protection policy overrides the backup host and storage unit information from the script with the backup host and storage unit information from Data Manager. Both the manual backups and scheduled backups of these Oracle databases are sent to this storage unit. To display the storage units and their assigned databases on the Oracle RMAN agent host, run the `ddutil -s` command.



Figure 7 Storage unit setup

Note: Oracle RMAN agent 19.6 and earlier releases do not support the mapping structure that allows protection policies to share the same storage unit. Backups of databases that are protected by older agents and different policies cannot target the same storage unit. Data Manager 19.7 and later releases contain logic that detects this condition when you add or edit a protection policy. The policy rules alert you to the conflict and fall back to the previous structure that mapped one policy to one storage unit. You can resolve this condition by upgrading the Oracle RMAN agent to release 19.7 or later.

1.14 Top-level directory changes

Only if the auto backup is enabled for the protected database and you have created a self-service protection policy for Oracle, complete the required top-level directory changes:

- Log in to the Oracle host as an Oracle user.
- To obtain the top-level directory information, run the following command:

```
/home/oracle/opt/dpsapps/rmanagent/ddutil -s
```

- To complete the changes to the control file configuration for the Oracle database, run the following RMAN command, which includes the top-level pathname from the `ddutil -s` command output:

```
CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE 'SBT_TAPE' TO './
< Top-Level Path>/%F';
```

After you run this command, all the database backup pieces including the auto backups will be written under the top-level directory created in the storage unit.

Note: This setting is required only if the backup is done using the self-service protection policy. For centralized protection policy, this setting is done automatically by Data Manager as seen in screenshot below.

```

...
RMAN> CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE 'SBT_TAPE' TO
'/DICTIS-159ab5df-da20-4137-8e65-e1161e75b5de/
...

```

Figure 8 Log of centralized protection policy showing automatic setup of the top-level directory



Oracle RMAN agent backup workflow

Oracle Database can be backed up using centralized protection policy and self-service protection policy. Based on the type of protection policy backup workflow changes. This section discusses workflow in each case.

1.15 Centralized protection backup

In the centralized protection backup, the entire backup life cycle is governed by Data Manager. There is no need for the DBA to create any RMAN scripts as all parameters are passed by Data Manager agent as per the backup options selected during protection policy creation.

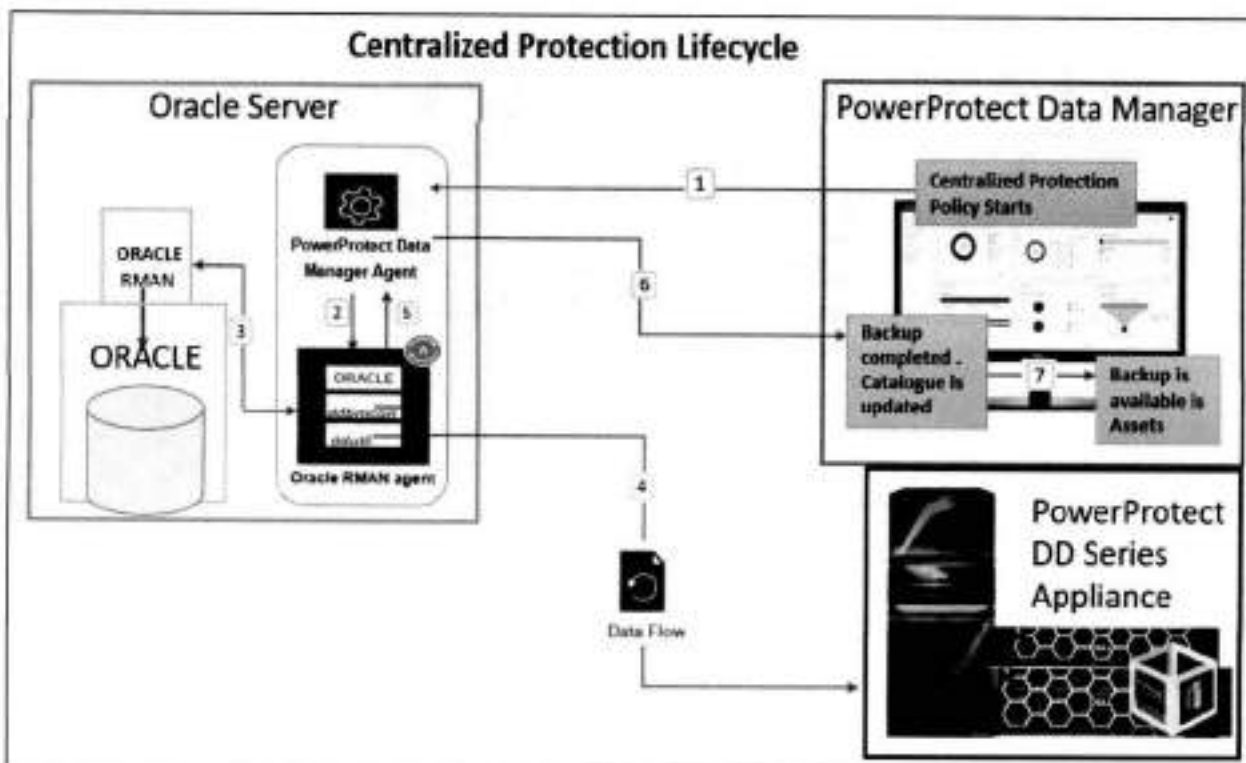


Figure 9 Centralized protection policy backup workflow

The backup workflow is as follows:

1. As scheduled, the centralized protection policy starts from Data Manager, and it communicates to the Agent service on Oracle client to start the backup process.
2. The agent service starts the Oracle RMAN agent and passes the information of assets that needs be backed up.
3. The Oracle RMAN agent connects to Oracle RMAN using the authentication provided, and Oracle RMAN starts Oracle databases backup and send it to Oracle RMAN agent.
4. The Oracle RMAN agent transfers the backup to DD series appliance, the number of streams used to transfer the data can vary based on the parallelism option selected for each asset.
5. When the data transfer is finished, Oracle RMAN agent sends the backup completion status to Agent service.
6. The agent service updates the backup status to Data Manager, and the backup catalog is updated with this information.

7. If the backup was successful, the backup copy is displayed in **Infrastructure > Assets > View Copies**.



Figure 10 Check backup copies



Figure 11 Checking available backup copies of Oracle host

1.15.1 Backup levels for centralized protection

The backup levels available during centralized protection policy are explained in table below:

Backup level	Description	Minimum frequency recommendation
Full	Backs up all the data.	Daily
Incremental Cumulative	Backs up only the data that has changed since the last full backup.	12 hours
Incremental Differential	Backs up 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.	6 hours

Log	Backs up the archived logs	30 minutes
-----	----------------------------	------------

You can define in what intervals these backups should run. Option to delete the logs after successful backup is also available under **Options->Delete archive log older than (Days)**.

Note: To delete the archived logs that are older than the specified number of days, ensure that the log backup option is enabled when you create the backup schedule. To delete the archived logs immediately after the log backups, set the flag option in `rman_agent.file`, available in the directory `$RMAN_AGENT_HOME/config` with the entry `DELETE_ARCHIVE_LOGS=TRUE`.

1.15.2 Enable multistream backup

To enable multistream Oracle backups for a centralized protection policy, you can set the parallelism value as the number of Oracle backup channels in the Data Manager UI. As an alternative, you can set the `PARALLELISM` parameter in the configuration file `rman_agent.cfg`.

Determine the required number of Oracle backup channels based on the system capacity. With the parallelism setting, you can override the number of backup channels from the Oracle RMAN agent client side. In the Data Manager UI, perform the following steps to set the parallelism for multistream backups:

1. Select **Infrastructure > Assets > Oracle**
2. Select the Oracle asset. Select **More Actions > Set Parallelism**.



1.15.3 Archive log backup

While setting up the centralized protection policy under **Options** page, select the settings for archive logs deletion from the production Oracle host using one of the following options:

- Do not delete:** Select this option to prevent the deletion of archived logs during backups. To delete the archived logs, the database administrator must run the delete command manually

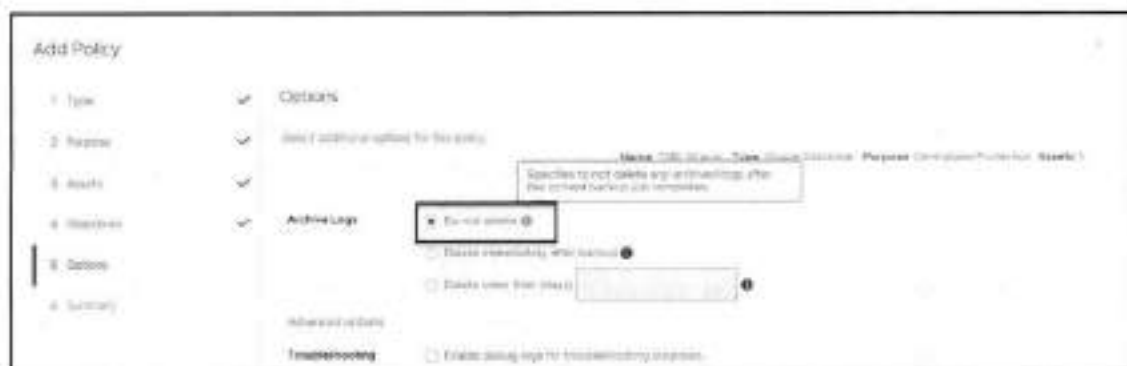


Figure 12 Archive log setting (1)

- Delete immediately after backup:** Select this option to enable the deletion of archived logs immediately after all the backup types that are performed through the protection policy



Figure 13 Archive log setting (2)

- Delete older than (days):** Select this option to enable the deletion of the available archived logs that are older than the specified number of days, for all the backup types that are performed through the protection policy. Set the number of days after which the archived logs are deleted.



Figure 14 Archive log setting (3)

1.15.4 Monitoring jobs and task for centralized protection policy

Use the **Protection Jobs** and **System Jobs** windows in the PowerProtect Data Manager UI to monitor the status of Oracle backup and to 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.

You can also view details for a job group and individual jobs and tasks. 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. This information enables you to monitor the status of individual jobs and tasks, view job and task details, and perform certain operations on jobs and tasks.

Use the **Group by** filter in the **Job ID Summary** window to view the application assets that are protected for all hosts in a protection job group. You can filter jobs by host for Microsoft SQL, Exchange databases, Oracle databases, File Systems, and SAP HANA databases.

To filter application assets by hostname, click the job ID for the job group, and then select **Group by > Host**. To display all assets in the job group, select **Group by > None**.



Figure 15 Monitoring backup using filter option Group by Host

Note: For Oracle hosts, the **Group by >Host** selection shows the job progress as 0% and then successful 100% once completed which means it will not show real-time progress. This is a known limitation with Oracle database backups. For more filter options and details, see the document [PowerProtect Data Manager Administration and User Guide](#).

1.16 Self-service protection backup

When the Data Manager admin creates a protection policy for Oracle databases, the self-service protection option enables the data owner to perform the manual backup operation from the command-line interface. In this option, the DBA creates their own RMAN scripts to back up the data directly to the DD series appliance using only the Oracle RMAN agent and Data Manager agent service to do the discovery every hour to check if the RMAN catalog is updated with any new backups.

To identify the storage unit and DD series appliance hostname, run the `ddutil -s` command on the Oracle client. Only if the autobackup is enabled for the protected database and you have created a self-service protection policy for Oracle, complete the required top-level directory changes as explained in [Top-level directory changes](#).

Following information is needed before performing the self-service protection of Oracle database. These parameters are required in the RMAN scripts:

- **SBT_LIBRARY**-The installation directory of the DD Boost library file - `libddobk.so`. The default installation directory is: `$RMAN_AGENT_HOME/lib`
- **STORAGE_UNIT**-The name of DD series appliance storage unit, which is created automatically when you add the protection policy. To display the storage units and their assigned databases on the Oracle RMAN agent host, run the `ddutil -s` command
- **BACKUP_HOST**-The hostname or IP address of the DD series appliance.
- **RMAN_AGENT_HOME**-The Oracle RMAN Agent software installation directory.

The following example shows an RMAN script that performs a full backup of the database and its archive logs:

```

connect target username/password:

run {
allocate channel c1 type SBT_TAPE parms 'SBT_LIBRARY=rman_agent_home/lib/
libddobk.so, ENV={RMAN_AGENT_HOME=rman_agent_home, STORAGE_UNIT=XYZ, BACKUP_HOST=bu-
ddbealin-17.lax.enc.com}';

backup database include current controlfile format 'AU' plus archivelog;

release channel c1;
}

```

Figure 16 Script for full backup and its archive logs with one channel

```

connect target username/password:

run {
allocate channel c1 type SBT_TAPE parms 'SBT_LIBRARY=rman_agent_home/lib/
libddobk.so, ENV={RMAN_AGENT_HOME=rman_agent_home, STORAGE_UNIT=XYZ, BACKUP_HOST=bu-
ddbealin-17.lax.enc.com}';
allocate channel c2 type SBT_TAPE parms 'SBT_LIBRARY=rman_agent_home/lib/
libddobk.so, ENV={RMAN_AGENT_HOME=rman_agent_home, STORAGE_UNIT=XYZ, BACKUP_HOST=bu-
ddbealin-17.lax.enc.com}';
allocate channel c3 type SBT_TAPE parms 'SBT_LIBRARY=rman_agent_home/lib/
libddobk.so, ENV={RMAN_AGENT_HOME=rman_agent_home, STORAGE_UNIT=XYZ, BACKUP_HOST=bu-
ddbealin-17.lax.enc.com}';
allocate channel c4 type SBT_TAPE parms 'SBT_LIBRARY=rman_agent_home/lib/
libddobk.so, ENV={RMAN_AGENT_HOME=rman_agent_home, STORAGE_UNIT=XYZ, BACKUP_HOST=bu-
ddbealin-17.lax.enc.com}';

backup database include current controlfile format 'AU' plus archivelog;

release channel c1;
release channel c2;
release channel c3;
release channel c4;
}

```

Figure 17 Allocated more channels for parallel backup using RMAN script

The backup workflow is as follows:

1. The DBA starts the Oracle RMAN script and script starts the Oracle database backup.
2. Oracle RMAN connects with Oracle RMAN agent to start the data transfer.
3. The Oracle RMAN agent opens connections with the DD series appliance and starts the data transfer.
4. When the backup is completed, the backup catalog information is passed from Oracle RMAN agent to Data manager agent service on client.
5. The PowerProtect Data Manager agent updates Data Manager with the backup status.
6. Data Manager updates the catalog information and sets the retention for oracle backup as per the self-service policy.
7. You can now see the backup copies under **Infrastructure > Assets > View Copies**.

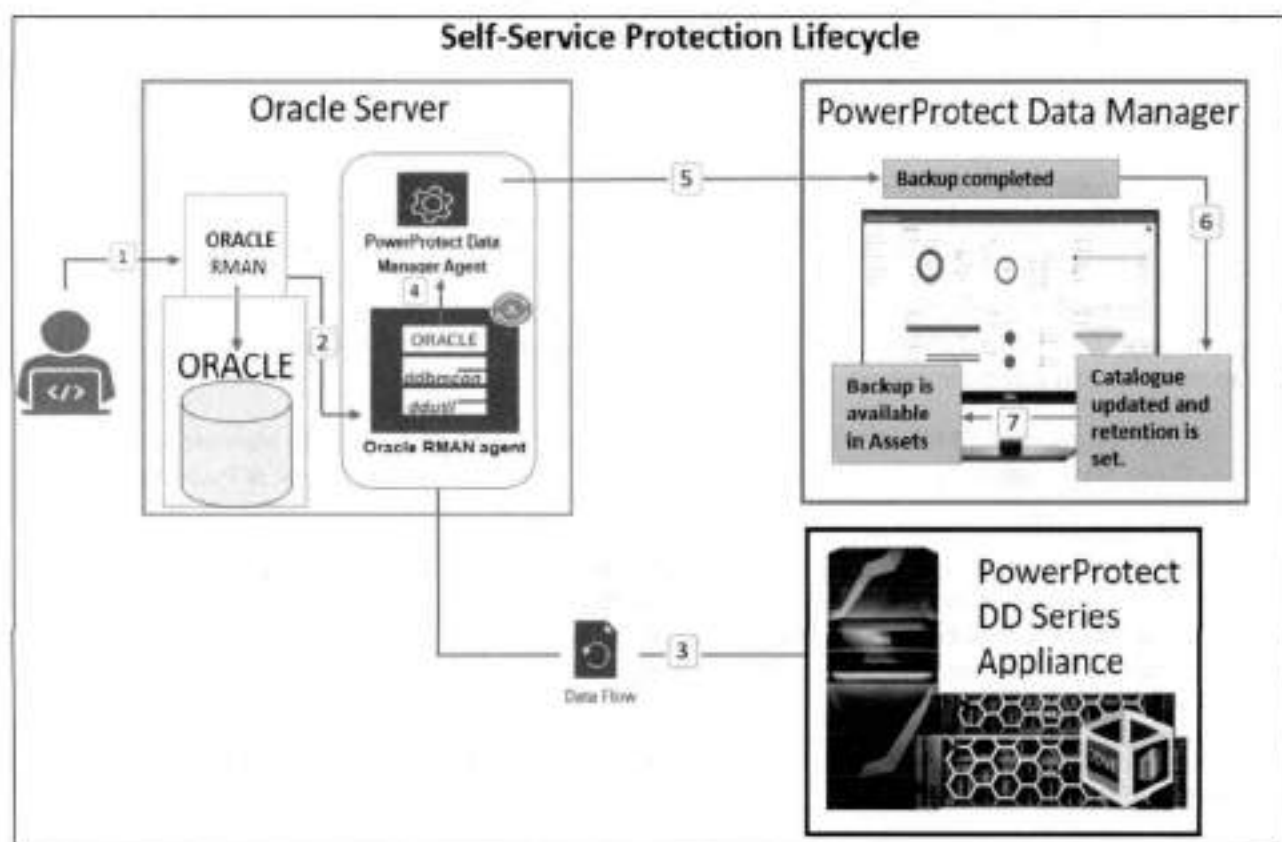


Figure 18 Self-service protection policy backup workflow

Oracle RMAN recovery

With Data Manager, you can perform centralized restore and recovery of Oracle backups from Data Manager UI or an Oracle database recovery on the Oracle database host by running one of the supported Oracle backup or recovery tools.

- Oracle Recovery Manager (RMAN) with `rman` command
- Oracle Enterprise Manager UI

1.17 Centralized restore and recovery of Oracle backups

Starting with Data Manager 19.11, you can perform centralized Oracle restore and recovery from the PowerProtect Data Manager UI when the Oracle Server data has been backed up through a protection policy. The centralized Oracle restore and recovery operations include the restore and recovery of a full online database, restore of only archive logs, and disaster recovery of an entire database.

You can perform the following types of centralized restore and recovery of Oracle Server backups:

- Centralized restore and recovery of a full online database without restore of the control file
- Centralized restore of only archive logs without restore of the control file
- Centralized disaster recovery of an entire database, including the spfile and control file

Using a centralized restore from the PowerProtect Data Manager UI, you can restore the Oracle database or archive logs for a single asset. You can perform the centralized restore to either the original Oracle Server host or an alternate host with the following requirements:

- The Oracle RMAN agent software must be installed and configured on the alternate host.
- The alternate host must be an Oracle server host that is a discovered asset of PowerProtect Data Manager.

Optionally, you can select to perform a dry run of any centralized restore and recovery operation to either the original host or an alternate host. The dry run option creates the required RMAN restore scripts but does not perform an actual restore or recovery. RMAN script will be created in the `$RMAN_AGENT_HOME/tmp` directory on the selected target host. You can use the RMAN restore scripts that the dry run creates to perform a self-service restore as required.

Centralized restore and recovery operations can be performed from the **Restore > Assets > Oracle** window in the PowerProtect Data Manager UI.



Figure 19 Centralized restore and recovery of Oracle backups

1.17.1 Centralized Oracle restore and recovery of a full online database

A centralized restore and recovery of an online Oracle database supports the following features:

- Restore and recovery to the original Oracle host.
- Restore and recovery to an alternate Oracle host.
- Point-in-time restore and recovery, based on time, SCN, or log sequence.
- Dry run of the database restore and recovery.

Note: A centralized restore and recovery of an Oracle database is allowed only when the backup copies location is listed as LOCAL or Local_Recalled in the Location column in the PowerProtect Data Manager UI. (To see the Location column, go to Infrastructure > Assets, select the asset on the Oracle tab, click View Copies, and then click the storage icon on the left.) To recall a cloud tier backup before you perform a centralized restore, follow the procedure "Restore the cloud tier backups to DD" mentioned in the

For any centralized restore to an alternate host, ensure that the alternate host is an Oracle Server host that is a discovered asset of PowerProtect Data Manager.

Note: If the alternate host is not included in the list of available hosts, follow the instructions to install and configure the Oracle RMAN agent on the alternate host. Ensure that the Oracle Server host is registered to the same PowerProtect Data Manager server.

You can use the PowerProtect Data Manager UI to perform a centralized restore and recovery of a full Oracle database backup without the control file.



Use the **Location** page to select the preferred type of restore.

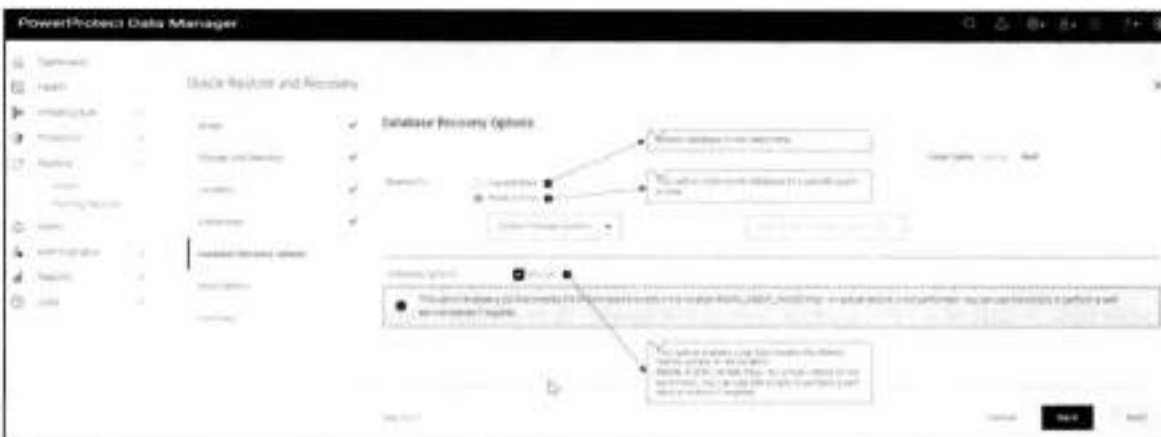
The option **Restore to original host** specifies to restore to the original host with the displayed hostname. If the original host is part of a RAC cluster, select the available node hostname from the list.



The **Restore to alternate host** option specifies to restore to an alternate host. Select the alternate hostname from the list, and then select the required instance name.



The **Database Recovery Options** page is shown below.



The **Restore To** setting enables you to select **Current time** or **Point in time**.

The **Current time** is the latest backup time in the control file.

For **Point in time**, select one of the following options from the menu:

- **System Change Number:** Enter the System Change Number (SCN) in the text box.
- **Timestamp:** Enter the date and time in the text box or click the icon to display a calendar and select the date and time.
- **Log Sequence:** Enter the log sequence in the text box.
- **Database Options:** Select **Dry Run** if you do not want to run an actual restore and recovery.

Use the **More options** page to specify the required options:

- **Set Stream Count:** Enter an integer stream count in the text box, if required. The default stream count is 4. The maximum stream count is 255
- **Compressed Restore:** To enable restore compression and reduce the impact on the network bandwidth, select **Use PowerProtect DD Boost (compressed restore)**.
- **Troubleshooting:** To enable troubleshooting, select **Enable debug log**.



1.17.2 Centralized Oracle restore of archive logs

A centralized restore of only Oracle archive logs supports the following features:

- Restore to the original Oracle host
- Restore to an alternate Oracle host
- Restore of a specific range of archive logs, based on time, SCN, or log sequence
- Dry run of the archive log restore

You can use the PowerProtect Data Manager UI to perform a centralized restore of an Oracle archive logs backup without the control file.

Use the **Scope** page to select online database recovery and to choose *Archive log only restore*.



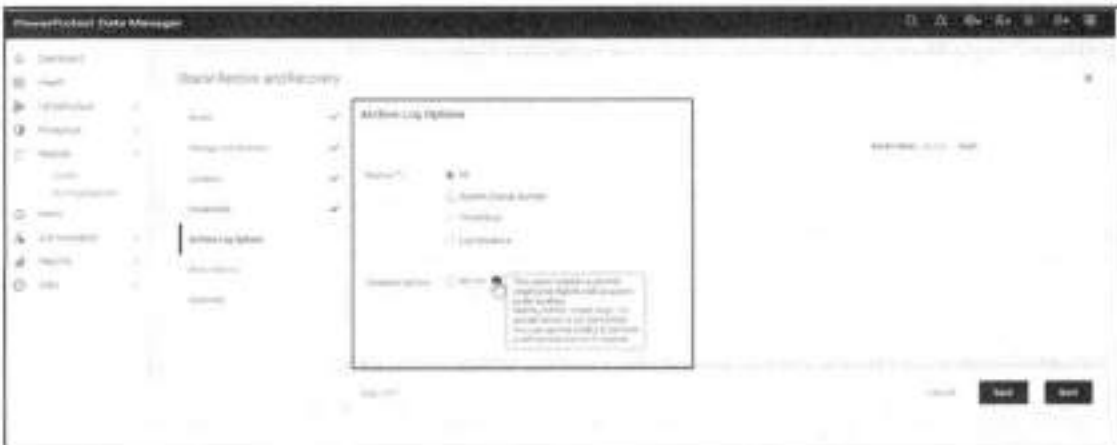
Use the **Location** page to select the preferred type of restore.



The **Restore to original host** setting enables restoring to the original host with the displayed hostname. If the original host is part of a RAC cluster, select the available node hostname from the list.

The **Restore to alternate host** setting specifies restoring to an alternate host. Select the alternate hostname from the list, and then select the required instance name.

Use the **Archive Log Options** page to select the required options.



Restore To: Select one of the following options:

- **All:** Specifies to restore all the archive logs of the database asset
- **System Change Number:** Specifies to restore the archive logs for the SCN range. Type the SCN start and end values in the text boxes
- **Timestamp:** Specifies to restore the archive logs for the timestamp range. Type or select the timestamp start and end values in the text boxes
- **Log Sequence:** Specifies to restore the archive logs for the log sequence range. Type the log sequence start and end values in the text boxes

Database options: Select Dry Run if you do not want to run an actual restore.

Use the **More Options** page to specify the required options:

- **Set Stream Count:** Enter an integer stream count in the text box, if required. The default stream count is 4. The maximum stream count is 255
- **Compressed Restore:** To enable restore compression and reduce the impact on the network bandwidth, select **Use PowerProtect DD Boost (compressed restore)**.
- **Troubleshooting:** To enable troubleshooting, select **Enable debug log**.



1.17.3 Centralized disaster recovery of an Oracle database

A centralized Oracle disaster recovery supports the following features:

- Restore to the original Oracle host
- Restore to an alternate Oracle host

Note: You can use the disaster recovery for Oracle testing and development purposes, for example, to validate the Oracle backups on an alternate host.

- Change of the database ID (DBID) of the restored database after disaster recovery
- Relocation of the Oracle data files
- Customization of Oracle startup parameter settings in the spfile of the Oracle database
- Point-in-time restore and recovery, based on time, SCN, or log sequence
- Dry run of the disaster recovery

You can use the PowerProtect Data Manager UI to perform a centralized disaster recovery of an Oracle database including the spfile and control file.

Use the **Scope** page to select Disaster recovery option.



This option restores and recovers the entire database, including the spfile and control file. You can use this option to restore the Oracle database to the original host or to an alternate host with a different database ID (DBID) than the original production host. The restore to the alternate host is usually used for test and development purposes.

Use the **Location** page to select the preferred type of restore.



The **Restore to original host** option specifies restoring to the original host with the displayed hostname. If the original host is part of a RAC cluster, select the available node hostname from the list.

The **Restore to alternate host** option specifies restoring to an alternate host. Select the alternate hostname from the list.

Use the **Folder Location** page to select one of the following options.



- Restore to original folder

Note: When you select to restore the Oracle data files to the original location, the original database is overwritten.

- Restore to alternate folder

To specify an alternate location, select Archive log, Control file, Datafiles, Fast recovery area, or Redo log from the menu, and then type the alternate location in the text box. For each additional alternate location, click the + icon, select the file type from the menu, and type the alternate location in the text box.

Use the **Instance Details** page to specify the required settings:



- **Oracle SID:** Enter the Oracle instance system ID (SID) in the text box.
- **Oracle Home:** Enter the valid Oracle home pathname in the text box.

Note: The Oracle home pathname must not include a final space or slash (/).

- **Credentials:** Select one of the following options:
 - Use the credentials set at asset level or policy level for restore.

Note: Credentials at the asset level take precedence over credentials at the protection policy level.

- Select the existing credentials or create new credentials for restore.

Use the **SPFILE Options** page to specify the required settings.



- **SPFILE Options:** Select this option to specify the restore of the spfile during the disaster recovery. This option is selected by default when the selected copy contains the spfile.

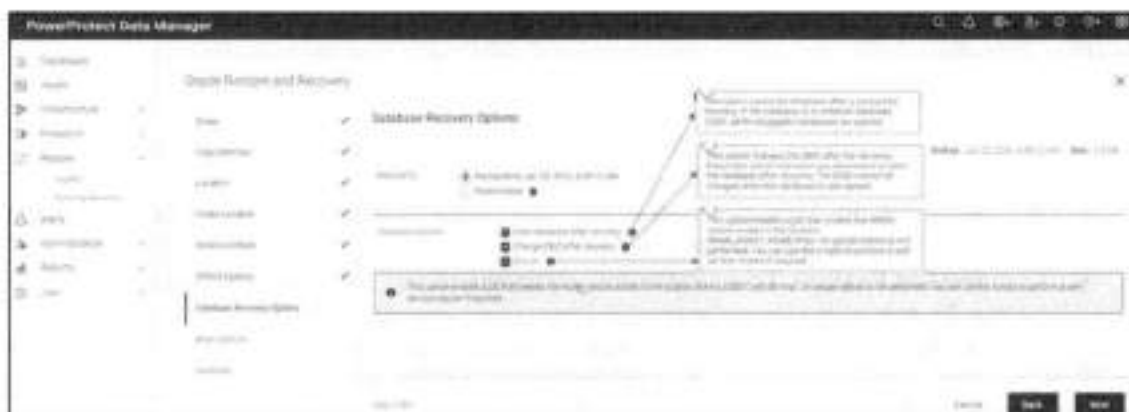
Note: If you do not select this option to restore the spfile, you must create the spfile manually and start the database instance in no mount mode.

- **Configure SPFILE Parameters:** To change a default spfile parameter setting, type the parameter name in the Parameter text box and the parameter value setting in the Value text box. For each

additional parameter setting that you want to change, click the + icon and then type the parameter name and value in the new blank text boxes.

For more details about configuring SPFILE parameters, see the document [PowerProtect Data Manager Oracle RMAN Agent User Guide](#) for detailed steps.

Use the **Database Recovery Options** page to select the required options.



- **Restore To:** Select Backup time (backup end time of selected backup copy) or Point in time.

For **Point in time**, select one of the following options from the menu:

- **System Change Number:** Type the System Change Number (SCN) in the text box.
 - **Timestamp:** Type the date and time in the text box or click the icon to display a calendar and select the date and time.
 - **Log Sequence:** Type the log sequence in the text box.
- **Database Options:** Select any required options from the list.
 - **Open database after recovery:** This option opens the database after the recovery. If the database is a container database (CDB), all the pluggable databases are opened.
 - **Change DBID after recovery:** This option changes the DBID after the recovery. Select this option only when you also select to open the database after the recovery. When the database is unopened, the DBID cannot be changed.

Note: When you select the Change DBID after recovery option, the asset is discovered automatically in PowerProtect Data Manager after a successful restore with the DBID change. When you do not select this option, the restored asset is not discovered automatically.

- Dry run

Use the **More Options** page to specify the required options.



- **Set Stream Count:** Type an integer stream count in the text box, if required. The default stream count is 4. The maximum stream count is 255.
- **Compressed Restore:** To enable restore compression and reduce the impact on the network bandwidth, select Use PowerProtect DD Boost compressed restore.
- **Troubleshooting:** To enable troubleshooting, select Enable debug log.

Centralized Oracle restore and recovery includes the following limitations:

- You cannot perform a cross-OS platform restore.
- You cannot perform a quick recovery.
- You cannot perform a centralized restore of an Oracle backup performed by a stand-alone Oracle RMAN agent before the agent was registered with PowerProtect Data Manager.

1.18 Self-service restores of Oracle databases

You can perform database restores directly to the Oracle application host by using the Oracle RMAN agent or you can restore an Oracle backup of a source client for disaster recovery or for a cross-restore to an alternate client host.

The restore workflow is as follows:

1. Oracle DBA starts the restore using RMAN recovery script.
2. Oracle RMAN connects to Oracle RMAN Agent and passes the information of assets to be recovered.
3. Oracle RMAN agents connects to DD series appliance and requests the data for recovery.
4. Data transfer starts from DD series appliance.
5. After recovery of database is completed, the script displays the completion status.




Figure 20 Self-service Oracle database restore workflow

To perform an Oracle database restore, you must prepare the database and then run an RMAN script to restore the data. The *RMAN documentation* provides detailed information about how to prepare the database and create the RMAN restore script. The documentation also describes all the supported restore features.

The following information is required before the Oracle database recovery operation:

- **SBT_LIBRARY:** The installation directory of the DD Boost library file - libddobk.so. The default installation directory is: \$RMAN_AGENT_HOME/lib
- **STORAGE_UNIT:** The name of DD series appliance storage unit, which is created automatically when you add the protection policy. To display the storage units and their assigned databases on the Oracle RMAN agent host, run the `ddutil -s` command
- **BACKUP_HOST:** The hostname or IP address of the DD series appliance.
- **RMAN_AGENT_HOME:** The Oracle RMAN Agent software installation directory.

To identify the storage unit and DD hostname, run the `ddutil -s` command on the Oracle client. For example, run the following command in the \$RMAN_AGENT_HOME/bin directory:

```
./ddutil -s
```

Specify the storage unit, top-level pathname, and DD hostname in the RMAN restore script.

The following example shows an RMAN script that performs a complete restore of the database to the current time, after the database has been prepared:

```
connect target username/password;

run {
set CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE 'SBT_TAPE' TO '._/
PLCTLP-4eb04bd9-b825-4e72-b668-14e9aacaa522/%F';

allocate channel cl type SBT_TAPE parms 'SBT_LIBRARY=rman_agent_home/lib/
libddobk.so, ENV=(RMAN_AGENT_HOME=rman_agent_home, STORAGE_UNIT=XYZ, BACKUP_HOST=bu-
ddbealin-17.iss.emc.com)';

restore database;
recover database;

release channel cl;
}
```

Figure 21 Restore script

Note: To increase the parallelism of the restore, you can allocate more channels.

For detailed steps to restore an Oracle application host or to restore to an alternate host, see [PowerProtect Data Manager Oracle RMAN Agent User Guide](#).

Oracle Data Guard support

Data Guard provides a comprehensive set of services that create, maintain, manage, and monitor one or more standby databases to enable production Oracle databases to survive disasters and data corruption. Data Guard maintains these standby databases as transitionally consistent copies of the production database.

Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data. It is common at large enterprise sites to deploy critical Oracle databases in High Availability mode or Data Guard mode.

1.19 Data Guard configuration

Data Guard configuration consists of one primary database and one or more standby databases. The standby databases will be always in sync with the primary database.



1.20 Data Guard: standalone mode support

Starting with version 19.12, Data Manager supports standalone mode assets that simplify the management of primary/standby node protection and recovery for users.

The following are Oracle RAC preferred node UI features:

- The Data Guard asset name is displayed in the format DBNAME(DBUNIQUE)
- The user can use a single instance of Data Manager to protect all Data Guard nodes
- Each database in a Data Guard configuration is considered a separate entity with its individual protection schedule and workflow
- No role or correlation is shown between databases in a Data Guard configuration
- This is supported on both Linux and AIX



1.21 Data Guard: the recovery catalog option

In the Data Guard environment, it is recommended that the user use a recovery catalog to manage the RMAN metadata for all physical databases, including primary and standby databases.

RMAN uses the recovery catalog as the single source of truth for the Data Guard environment. RMAN does not automatically resynchronize every database in the Data Guard environment when connected as TARGET to one database in the environment.

This option is applicable only for centralized protection in an Oracle Data Guard environment. This option resynchronizes the recovery catalog either synchronously or asynchronously with backup copies after each backup.



1.22 Data Guard: the disaster recovery option

The disaster recovery option helps to restore the database as either a primary or standby database in a Data Guard configuration to a specified point in time. The user can bring up a standby database using a primary or standby database backup, or bring up a primary database using a primary or standby database backup (by using spfile/pfile).



Note: After a disaster recovery, have an administrator get the Data Guard in sync.

1.23 Data Guard: self-service protection

For self-service protection, synchronize the recovery catalog in separate RMAN sessions at regular intervals to update the recovery catalog with the current metadata from the target database control file.

The user can run the `resync catalog` command to initiate a full resynchronization of the recovery catalog.

To enable self-service backups for an Oracle RAC Data Guard configuration, add the `DB_UNIQUE_NAME` parameter to the RMAN script `ALLOCATE CHANNEL`.

Replication and DD Cloud Tier

During the protection policy creation self-service or centralized, you can add the replication to a remote PowerProtect DD series appliance as the replication target.

In a protection policy, click **Replicate** next to **Primary Backup**, **Primary Retention**, or **Extend Retention**. An entry for **Replicate** is created to the right of the primary or extended retention backup schedule. Under **Replicate**, click **Add**. The **Add Replication** dialog appears.

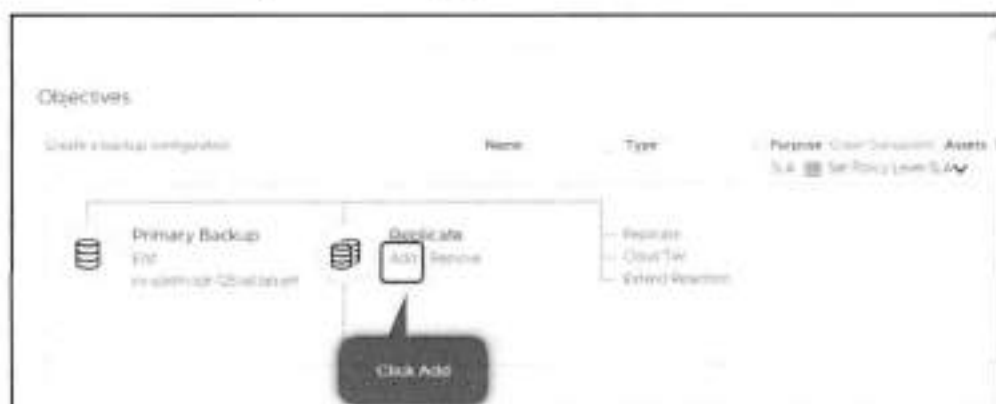


Figure 22 Replication configuration

Complete the schedule details in the **Add Replication** dialog and click **Save** to save your changes.



Figure 23 Replication schedule

Data Manager cloud tier feature works in tandem with the DD Cloud Tier feature to move Data Manager backups from DD series appliance to the cloud. This provides long-term storage of Data Manager backups by seamlessly and securely tiering data to the cloud. From the Data Manager UI, you configure cloud tier to move Data Manager backups from DD series appliance to the cloud, and you can perform seamless recovery of these backups. DD series appliance cloud storage units must be preconfigured on the DD series appliance before they are configured for cloud tier in the Data Manager UI. The [PowerProtect Data Manager Administration Guide](#) provides more information.

Both Oracle centralized and self-service protection policies support cloud tiering. You can create the cloud tier schedule from both primary and replication stages. Schedules must have a minimum weekly recurrence and a retention time of 14 days or greater. Ensure that the DD series appliance is set up for cloud tiering and follow the below step:

1. Click **Cloud Tier** next to **Primary Backup** or **Extend Retention**, or if adding a cloud stage for a replication schedule that you have added, click **Cloud Tier** under **Replicate**. An entry for **Cloud Tier** is created to the right of the primary or extended retention backup schedule, or below the replication schedule.

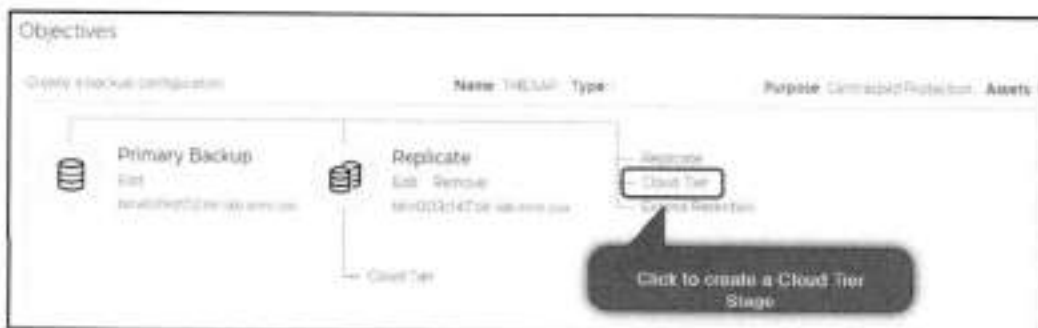


Figure 24 Cloud Tier Configuration

2. Under the entry for **Cloud Tier**, click **Add**.



Figure 25 Cloud Tier Configuration

3. The **Add Cloud Tier Backup** dialog appears, with summary schedule information for the parent node to indicate whether you are adding this cloud tier stage for the primary backup schedule, the extended retention backup schedule, or the replication schedule.

Complete the schedule details in the **Add Cloud Tier Backup** dialog, and click **Save** to save your changes.

Figure 26 Cloud Tier Configuration

4. The Protection Policy Summary Lists **Replicate** and **Cloud Tier**.



Figure 27 Replication and DD Cloud Tier

A Technical support and resources

Dell.com/support is focused on meeting customer needs with proven services and support.

The [Data Protection Info Hub](#) provides expertise that helps to ensure customer success with Dell data protection products.

A.1 Related resources

- [PowerProtect Data Manager for Oracle RMAN agent User Guide](#)
- [PowerProtect Data Manager Administration and User Guide](#)
- [PowerProtect Data Manager Deployment Guide](#)
- [DDOS Administration Guide](#)

PowerProtect Data Manager E-LAB Navigator: Provides compatibility information, including specific software and hardware configurations that PowerProtect Data Manager supports. To access E-LAB Navigator, go to [PowerProtect Data Manager Compatibility Matrix](#).

Dell EMC PowerProtect DD Virtual Edition on Amazon Web Services

December 2021

H18632.1

White Paper

Abstract

This white paper explains the steps to deploy Dell EMC PowerProtect DD Virtual Edition (DDVE) on Amazon Web Services (AWS).

Dell Technologies

Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA December 2021 H18632.1.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Contents

Executive summary	4
PowerProtect DD Virtual Edition	4
Deploying DDVE on AWS	7
Configuring DDVE on AWS	20
Best practices	24
Conclusion	26
Technical support and resources	26

Chapter 1	Deploying DDVE on AWS
Chapter 2	Configuring DDVE on AWS
Chapter 3	Best practices
Chapter 4	Conclusion
Chapter 5	Technical support and resources

Executive summary

Overview With more organizations moving to cloud platforms, one of the biggest concerns is how to manage cloud data protection. Dell Technologies provides a cloud-enabled data protection solution that manages long-term retention to the cloud with cost-effective solutions.

Dell EMC PowerProtect DD Virtual Edition (DDVE) is a software-defined data protection solution that brings efficient and reliable data protection to remote and branch office, entry-level, and cloud environments. This white paper discusses the prerequisites, and how to deploy and configure PowerProtect DDVE on Amazon Web Services (AWS).

Audience This white paper is intended for Dell Technologies customers, partners, and employees looking for options to protect the workloads hosted on Google Cloud Platform using PowerProtect DDVE.

Revisions

Date	Description
January 2021	Initial release
December 2021	Updated for PowerProtect DD Virtual Edition 7.7 release; template update

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#) (subject line: Feedback for document: H18632.1).

Author: Charu

Note: For links to other documentation for this topic, see the [Data Protection Info Hub](#).

PowerProtect DD Virtual Edition

Introduction to DDVE

DDVE is a software-defined data protection solution with the PowerProtect DD series appliance, including all the core differentiating features of the DD series.

DDVE runs the DD Operating System (DDOS) and includes the DD System Manager user interface (UI) and the DDOS command-line interface (CLI) for performing system operations.

DDVE includes the following features:

- High-speed, variable-length deduplication for a 10 to 30 times reduction in storage requirements
- Unparalleled data integrity to ensure reliable recovery, and seamless integration with leading backup and archiving applications
- DD Boost to speed backups by 50 percent

- DD Encryption for enhanced security of data
- DD Replicator for network-efficient replication that enables faster time-to-DR readiness

DDVE can be deployed on any standard hardware, converged or hyperconverged, and runs in VMware vSphere, Microsoft Hyper-V, KVM. It also runs in-cloud with Amazon Web Services (AWS) (cloud and gov cloud), VMware Cloud (VMC), Azure (cloud and gov cloud), and Google Cloud Platform (GCP). DDVE is also certified with Dell EMC VxRail and Dell EMC PowerEdge servers.

DDVE scales up to 256 TB (in-cloud AWS, Azure, and Google Cloud) and up to 96 TB (on-premises) per instance.

DDVE cloud features

DDVE provides the capabilities of a cloud DD system using the following resource configuration sizes:

- DDVE on Block storage up to 16 TB
- DDVE on S3 storage up to 256 TB

DDVE features supported in AWS:

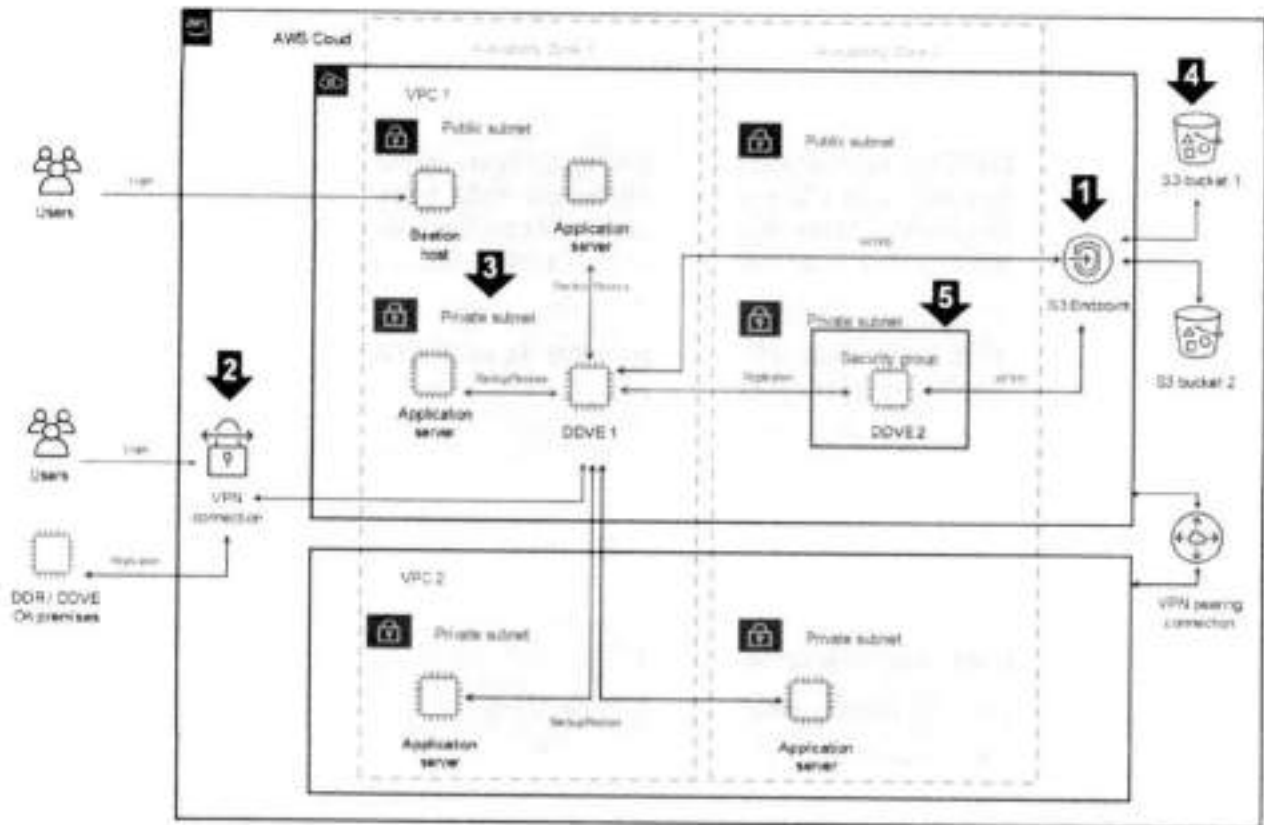
- DD Boost managed file replication (MFR)
- Encryption
- MTree replication
- DD System Manager UI for DDVE management
- DD Active Tier (DD Cloud Tier is not supported)
- Secure multitenancy (SMT) with Network Isolation Support
- DD Boost/BoostFS for big data
- Key Management Interoperability Protocol (KMIP)
- More restricted IPTables settings
- AWS for Government Cloud
- Retention Lock Governance Edition (supported on DDVE, on premises and in the cloud)
- Auto Retention Lock (ARL) and Indefinite Retention Hold (IRH), supported for RLG MTrees on DDVE

DDVE replication capabilities:

- Managed file replication and MTree replication
- Replication across availability zones and regions
- Bi-directional replication between on-premises and AWS

DDVE on AWS architecture

The following figure shows the architecture of DDVE on AWS:



1. To keep data traffic between DDVE and the S3 bucket within the AWS infrastructure, it is recommended to create an S3 endpoint. The S3 endpoint keeps DDVE from depending on a NAT Gateway or Public IP address to access the S3 bucket.
2. To keep data transfers secure, it is recommended to use a VPN connection to replicate data from an on-premises host to DDVE in the cloud or the opposite way.
3. DDVE is categorized as a backend server. It must be kept in a private subnet with a private address. Never set a public IP address for DDVE.
4. It is recommended to create the S3 bucket in the region where the DDVE instance is running. A separate bucket per each DDVE is required.
5. All DDVE instances must be secured with the appropriate security group entries.

Notes:

- Typically, SSH (Port 22) or HTTPS (Port 443) is used for DDVE inbound access.
- HTTPS (443) must be allowed for outbound S3 bucket access for DDVE.
- TCP ports 2049 and 2051 are used for DD Boost and replication purposes.

Deploying DDVE on AWS

Prerequisites to deploy DDVE on AWS

The general requirements for deploying DDVE on AWS are as follows:

- **Create an AWS account:** To deploy DDVE on AWS, an AWS account must be created.
- **Identity and access management:** AWS recommends that IAM user or role for authenticating with AWS must be created and root credentials should never be used to deploy the CloudFormation template. The IAM user must be allowed to perform AWS CloudFormation actions. The EC2 instance must be granted the IAM role to provide permissions to S3 storage.

Preparing environment to deploy DDVE on AWS

DDVE running on AWS cloud allows the customer to backup and restore the operational data from S3 object store.

The high-level steps involved are as follows:

1. Configure the network environment. For secure access to DDVE, it is recommended to use VPC architecture that AWS provides. Configure the following components:
 - VPC
 - Subnet
 - Route tables
 - Security groups
 - Network access control list
 - VPC Gateway endpoint for connectivity to S3
2. Create an S3 bucket.
3. Configure role-based access to the AWS object store.
4. For secure login to DDVE, create an EC2 key access pair.

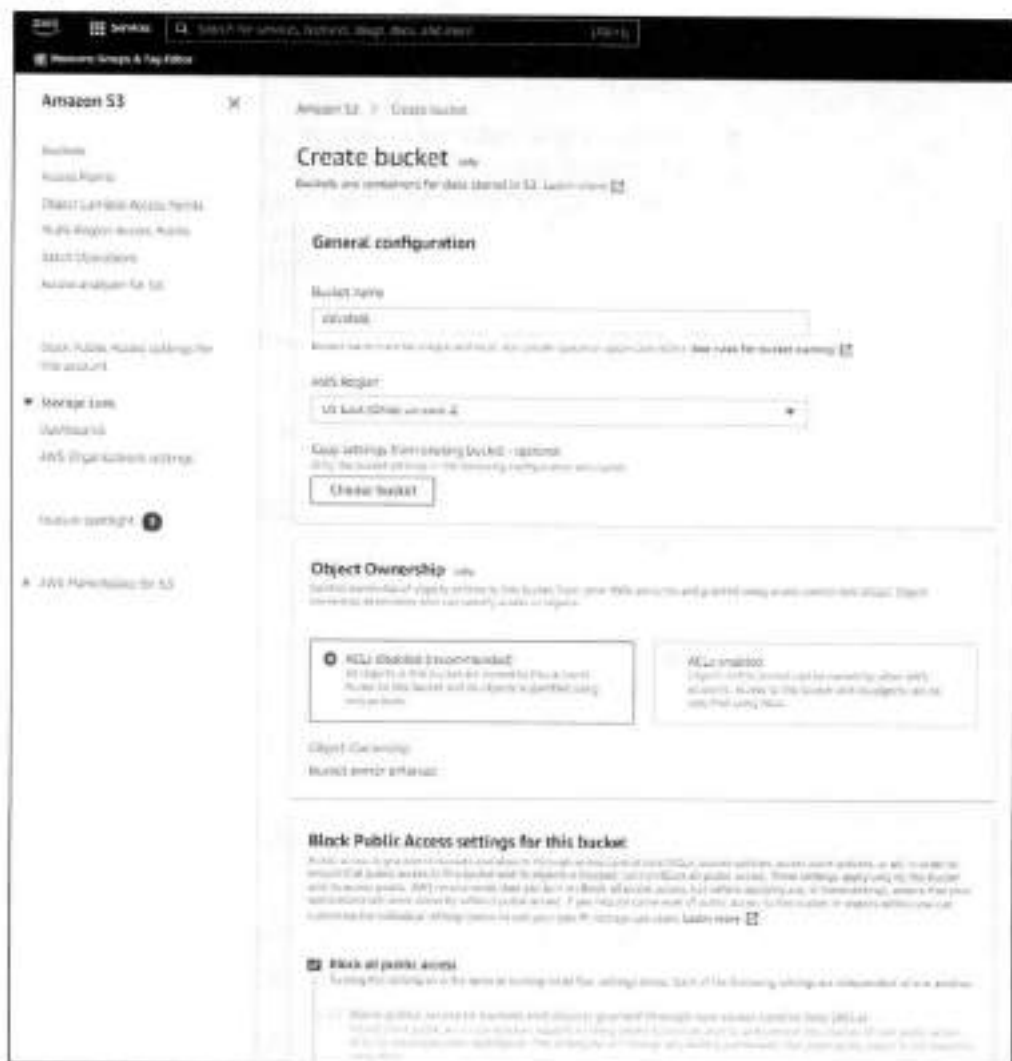
Creating an S3 bucket

Create a bucket in S3 and make note of the bucket name as it is used in further steps. The steps involved to create a bucket are as follows:

1. Log in to the AWS console, and select **Services > S3**.
2. Click **Create bucket**, and then enter the bucket name and region.
 - To access an S3 bucket, AWS recommends using hosted-style URLs (where domain name includes the bucket name) instead of path-style URLs. For hosted-style URLs to work, do not use dots (".") in the bucket name.
 - Create the bucket in the same region as the DDVE instance.
 - Provide a bucket name that is no longer than 48 characters.
 - Do not enable bucket versioning for the bucket that is associated with the DDVE. Versioning adds to storage costs because older versions of the objects

are retained despite running the DDVE garbage collection process. Enabling versioning can also cause potential performance issues.

3. Click Create Bucket.



Setting up a role-based access to the AWS object store

The object store in AWS uses role-based access for S3 access. To access the S3 bucket, create and attach the Identity and Access Management (IAM) role to DDVE.

Prerequisites: To create the IAM role and the policy that is associated with the role, the AWS user must have the necessary IAM privileges. The following IAM privileges and actions are required to create and attach the IAM role:

```


"iam:AddRoleToInstanceProfile",
  "iam:AttachRolePolicy",
  "iam:CreateRole",
  "iam>DeleteRole",
  "iam>DeleteRolePolicy",
  "iam:DetachRolePolicy",
  "iam:GetRole",
  "iam:GetRolePolicy",
  "iam:ListRolePolicies",
  "iam:ListRoles",
  "iam:PassRole",
  "iam:RemoveRoleFromInstanceProfile",
  "iam:UpdateRolePolicy",
  "iam:CreateInstanceProfile",
  "iam:PutRolePolicy",
  "iam>DeleteInstanceProfile"

```

When the role is attached to DDVE, the S3 object store credentials are automatically fetched. The AWS infrastructure periodically rotates the access credentials. The DDVE automatically fetches the new credentials before the old credentials expire.

Use the following procedure to set up role-based access to the object store:

1. Create the policy to attach with the IAM role:
 - a. Sign into the AWS Management Console and open the IAM Service Console.
 - b. In the navigation pane of the IAM console, select **Policies > Create policy**.
 - c. Create a policy for AWS Standard Cloud or AWS Gov Cloud:
 - d. In the **Create policy** web page, select the **JSON** tab. Replace the text under the JSON tab with the following content.



```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowS3Access",
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
        "s3:GetObject",
        "s3:ListBucket",
        "s3:DeleteObject"
      ],
      "Resource": [
        "arn:aws:s3:::*",
        "arn:aws:s3:bucket*"
      ]
    }
  ]
}

```

- e. Verify this information and click **Review policy**.
 - f. Provide a name and description for the policy and click **Create policy**.
2. Create the role for S3 bucket access:
 - a. In the navigation pane of the IAM console, select **Roles > Create role**.
 - b. On the **Create role** page:
 - Select **AWS service** to choose the type of trusted entity.

- Select **EC2** to choose the service that will use this role, and then click **Next Permissions**.
3. On the **Attach permissions policies** page, select the policy that is created in the previous section. Select **Next Tags** to create a tag for the role.
 4. Click **Next: Review**. In the **Review** section, provide a name for the role and click **Create role**.

Create role 1 2 3 4

Review

Provide the required information below and review this role before you create it.

Role name
Use alphanumeric and hyphen characters. Maximum 63 characters.

Role description
Maximum 1024 characters. Use alphanumeric and hyphen characters.

Trusted entities

Policies

Permissions boundary

No tags were added.

Attach the role to the DDVE instance before it can be configured. This task can be done during or after deployment.

Deploying DDVE on AWS

You can use the following methods to deploy DDVE on AWS:

- Cloud Formation Template (CFT)
- Amazon Machine Image (AMI)

CFT is considered as the recommended method to deploy DDVE on AWS as NVRAM and metadata disks are created and attached automatically in the correct order.

Deploying DDVE using Cloud Formation Template (CFT)

1. Click the appropriate link:
 - To deploy in AWS standard cloud, use <https://aws.amazon.com/marketplace>.
 - To deploy in AWS gov cloud, use <https://aws.amazon.com/mp/govcloud/>.

2. Search PowerProtect DD Virtual Edition.

The screenshot shows the AWS Marketplace search results for 'PowerProtect DD Virtual Edition'. The search bar at the top contains the text 'PowerProtect DD Virtual Edition'. On the left, there is a 'Refine results' sidebar with filters for Categories (Infrastructure Software (2), DevOps (1)), Delivery methods (CloudFormation Template (3), Amazon Machine Image (2)), and Publisher (Dell Technologies (3)). The main results area shows two items:

- Dell EMC PowerProtect DD Virtual Edition (DDVE)** by Dell EMC. Description: Dell EMC PowerProtect DD Virtual Edition (DDVE) is the software-defined data protection solution based on industry-leading Dell EMC PowerProtect DD, the world's most trusted protection storage. DDVE can now deliver increased transactional and operational efficiencies, reliability and lower TCO by...
- Dell EMC PowerProtect Data Manager and PowerProtect DD Virtual Edition** by Dell Technologies. Description: Data owners and administrators can deploy PowerProtect Data Manager and PowerProtect DD Virtual Edition in AWS to protect business-critical workloads in the cloud. PowerProtect Data Manager enables protection of traditional workloads including Oracle, SQL, SAP HANA and file systems as well as...

3. Select Dell EMC PowerProtect DD Virtual Edition (DDVE) and click Continue to Subscribe.

The screenshot shows the product page for 'Dell EMC PowerProtect DD Virtual Edition (DDVE)'. The page includes the Dell Technologies logo and a 'Continue to Subscribe' button. The product title is 'Dell EMC PowerProtect DD Virtual Edition (DDVE)'. Below the title, there is a 'Product Overview' section with the following text:

Dell EMC PowerProtect DD Virtual Edition (DDVE) is the software-defined data protection solution based on industry-leading Dell EMC PowerProtect DD, the world's most trusted protection storage. DDVE can now deliver increased transactional and operational efficiencies, reliability and lower TCO by offloading agent storage to Amazon S3. DDVE is now up to 20TB instances in-cloud and on-prem private supports AWS IoT Core.

Key features listed in the 'Highlights' section include:

- With the latest release of PowerProtect DD Virtual Edition, DDVE now supports 10 instances in AWS which expands the highest coverage including for Cloud End.
- Action increased transactional and operational efficiencies, reliability and lower TCO by offloading agent storage to Amazon S3.
- DDVE supports 2TB - 20TB on instances with the ability to expand capacity to 175 terabytes.

Technical specifications are listed in a table:

Version	PowerProtect DD Virtual Edition 7.7.0.0
By	Dell EMC, LP
OS	Linux
Supported Systems	CloudFormation, AWS, S3, S3 Glacier
Delivery Methods	Amazon Machine Image, CloudFormation Template

4. Click **Continue to Configuration**.

Dell Technologies Dell EMC PowerProtect DD Virtual Edition (DDVE) Continue to Configuration

Product Detail **Subscribe**

Subscribe to this software

You've subscribed to this software. Please see the terms and pricing details below or click the button above to configure your software.

Terms and Conditions

Dell EMC Offer

You have subscribed to this software and agreed that your use of this software is subject to the pricing terms and the seller's End User License Agreement (EULA). You agree that AWS may share information about the transaction (including your payment terms) with the respective seller, supplier or underlying provider, as applicable, in accordance with the AWS Privacy Notice. Your use of AWS services remains subject to the AWS Customer Agreement or other agreement with AWS governing your use of such services.

Product	Effective date	Expiration date	Action
Dell EMC PowerProtect DD Virtual Edition (DDVE)	10/10/2019	N/A	View Details

5. Select the following configuration, and then click **Continue to Launch**.

- **Fulfillment Option:** Select **Cloud Formation Template**.
- **Software Version:** Select the correct version.
- **Region:** Select where the DDVE is to deploy.

Dell Technologies Dell EMC PowerProtect DD Virtual Edition (DDVE) Continue to Launch

Product Detail **Configure**

Configure this software

Choose a fulfillment option and software version to launch this software.

Fulfillment option

CloudFormation Template

DDVE Deployment - CloudFormation

Software version

PowerProtect DD Virtual Edition 7.7.0.0 (Sep 24, 2)

Region

US East (Ohio)

Pricing information

This is an estimate of typical software and infrastructure costs based on your configuration. Your actual charges for such statement will vary after launch. For more information, see the pricing page.

Software Pricing

Dell EMC PowerProtect DD Virtual Edition (DDVE) **\$0/yr**

BYOL

- Review the configuration details, select **Launch the Cloud Formation template**, and click **Launch**. The template URL is populated.

Dell Technologies Dell EMC PowerProtect DD Virtual Edition (DDVE)

< Product Details | Subscription | Configure | **Launch**

Launch this software

Review the launch configuration details and follow the instructions to launch this software.

Configuration details

Fulfillment option	DDVE Deployment - CloudFormation (Recommended) Dell EMC PowerProtect DD Virtual Edition (DDVE)
Software version	PowerProtect DD Virtual Edition 7.7.0.0
Region	US East (Ohio)

[Usage instructions](#)

Choose Action

Launch CloudFormation

Choose this action to launch your configuration through the AWS CloudFormation console.

- Click **Next**.

CloudFormation | Stack | **Create stack**

Step 1: **Specify template**

Step 2: Specify stack details

Step 3: Configure stack options

Step 4: Review

Create stack

Prerequisite - Prepare template

Prepare template
First, create a template. A template is a JSON or YAML file that contains configuration information about the AWS resources you want to create in the stack.

Template is ready Use a sample template Create template in Designer

Specify template
A template is a JSON or YAML file that describes your stack's resources and properties.

Template source
Select a template source to provide a URL where it will be stored.

Amazon S3 URL Upload a template file

Amazon S3 URL

<https://s3.amazonaws.com/awstesting-fulfillment-of-templates-pod/1000/121-1001-4766-4052-11f96d815c73d006011-2204-4912-9600/1011001-Amazon-DDVirtualEdition-1001>

SS URL: <https://s3.amazonaws.com/awstesting-fulfillment-of-templates-pod/1000/121-1001-4766-4052-11f96d815c73d006011-2204-4912-9600/1011001-Amazon-DDVirtualEdition-1001>

- Enter the following values to create the stack:
 - Stack name
 - DDVE Model: Choose the appropriate model from the available options.

- Default number of metadata disks can be overridden by selecting a value from 1-24. The maximum number of metadata disks that can be attached to a DDVE instance in AWS is 24.

Note: Only the number of metadata disks can be overridden. The size of individual disks cannot be changed.

- DDVE name tag
- IAM Role for S3 access—Type the IAM role that was created while preparing the environment.
- Key pair—Select an existing key pair from the drop-down list.
- Subnet ID
- Security Groups

The screenshot shows the 'Specify stack details' step in the AWS CloudFormation console. The 'Stack name' is 'ddve-stack'. Under 'Parameters', the following values are entered:

- DDVE Instance Configuration:**
 - DDVE Model: Dell EMC DDVE
 - Metadata Disks: 24
- DDVE Name Tag:** ddve-stack
- IAM Role for S3 access:** ddve-s3-role
- Key Pair:** ddve-key-pair
- DDVE Network:** ddve-vpc
- Security Groups:** ddve-sg

At the bottom right, there are 'Cancel', 'Previous', and 'Next' buttons. A 'Next' button is highlighted.

9. Continue stack configuration as needed. Click **Next**.

10. Review the stack configuration and click **Create Stack**.

11. Check the status of the stack created.



12. When the stack creation is complete, go to the EC2 instances and select the region to deploy the DDVE. Use the DDVE name tag from step 8 and verify that the corresponding EC2 instance is running.

Note: Avoid disabling or modifying the primary interface settings. The primary interface in cloud deployments has the default gateway setting and is the only interface with which the DDVE can connect to the metadata server. The metadata server is critical for DDVE operation.




Deploying DDVE manually using Amazon Machine Image (AMI)

1. Sign into AWS management console and click on AWS marketplace.
2. Search for **PowerProtect DD Virtual Edition**.
3. Select **Dell EMC PowerProtect DD Virtual Edition (DDVE)** to choose the Amazon Machine Image (AMI).



- Choose the instance type from the supported instance types and admin guide can be referred for instance configuration details. Click **Configure Instance Details**.



- Select the VPC and subnet in which DDVE is to be deployed and select the IAM role that was created while preparing the environment. Click **Add Storage**.

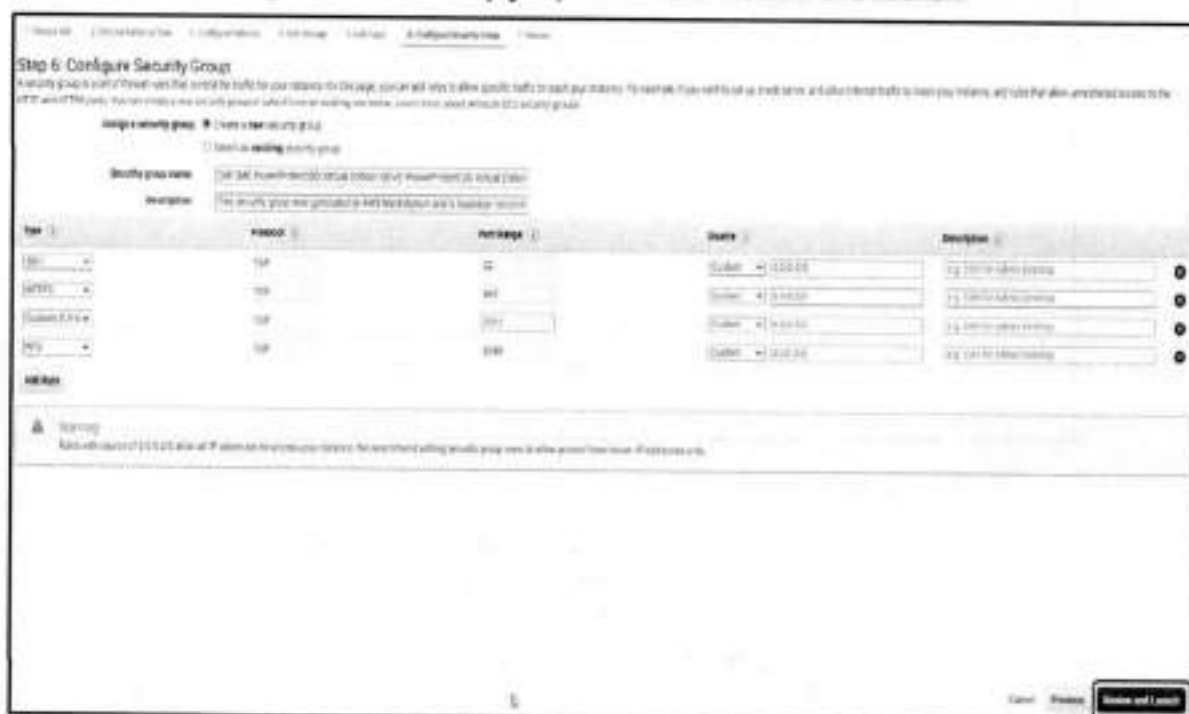


6. Add the NVRAM disk and metadata disks as shown below. Click **Add Tags**.**Notes:**

- It is important to add NVRAM disk before adding the metadata disks. Adding them in different order causes an unsupported hardware configuration error. Also, ensure that EBS volume type is GP2 for all disks.
- If required more metadata disks can be added to the DDVE instance from the AWS console. AWS DDVE instances support adding only up to 24 metadata disks. If the limit is reached for adding metadata disks, existing metadata disks can be expanded. For details about adding more metadata disks and expand metadata storage, see the document [Dell EMC PowerProtect DDVE on Amazon Web Services Installation and Administration Guide](#)

7. Adding tags is an optional step. Click **Configure Security Group**.

8. Create a new Security group or select from an existing one. See the administrator's guide for the security group details. Click **Review and Launch**.



9. Review the configuration details and click **Launch**.



10. Select a key pair value or create a new key pair value for this instance, and then click **Launch Instance**.



11. Click **View instances** to go to the EC2 instance tab.



12. The DDVE instance has been deployed successfully, and it is ready for configuration.



13. Click the instance to view the instance details. It provides information about the public and private IP address that can be used to configure the instance later.



Configuring DDVE on AWS

DDVE configuration can be done in two ways: DD system manager (UI) or command-line interface (CLI). In this white paper, configuring DDVE using DD system manager is discussed in detail.

1. Log in to DD System Manager using the DDVE IP address. The default login credentials for the DDVE instance are as follows:
 - **Username:** sysadmin
 - **AWS default password:** Default sysadmin password is the EC2 instance-id for the DDVE



2. A prompt window to update the password is shown. Change the password and click **Save**.



3. Accept the End User License Agreement.



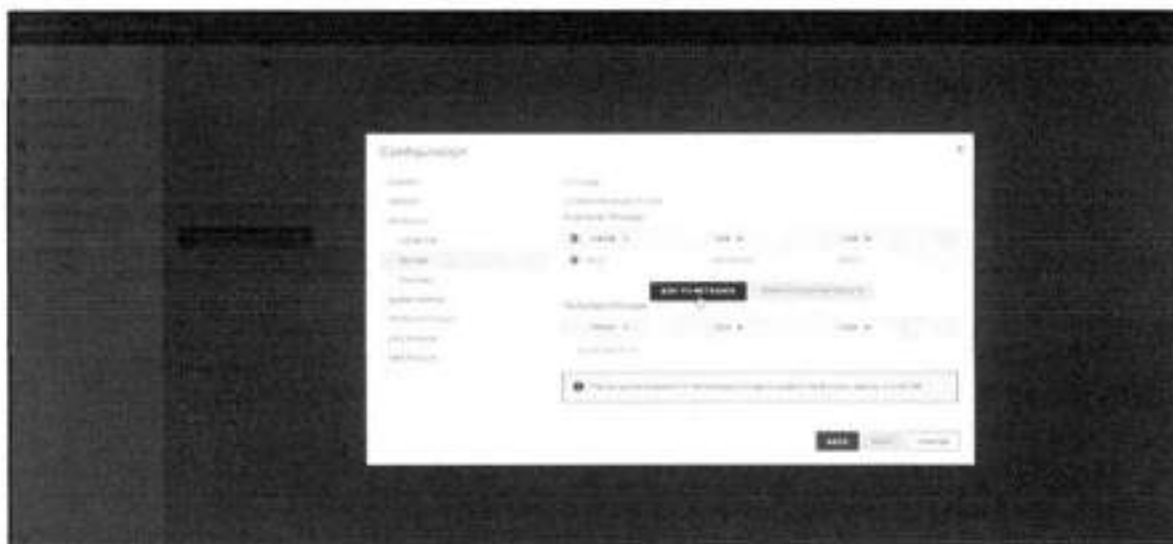
4. The configuration wizard is launched automatically. Configure the license if required and click **Next**.



5. Review the Network settings and select **NEXT** to proceed.
6. Click **Yes** to set up File System configuration.
7. For the **Storage Type**, select **Object Store** and enter the passphrase and the bucket name.



8. To configure the storage, under **Available Storage**, select the disks and click **Add to Metadata** to move them to the Metadata Storage section. Add the disks to the active tier to add the metadata storage disk to the instance.



9. On the **File System Summary Page**, select the **Summary** tab to review all the fields. Select **Enable file system after creation** and click **Submit**.

The file system is created and enabled:



10. Click **OK** to go to the **System Settings** tab.
11. In the system settings tab, passwords can be reset, and the email server can be configured if required.
12. Click **Submit** to save the system settings. Close the wizard.

Best practices

13. The summary of the file system can be seen under **Data Management > File System**.



14. DDVE must have accurate and consistent time synchronization for object store communication. DDVE can synchronize time by using Amazon Time Sync Service or by configuring an NTP server.

Note: If any changes are required after the initial DDVE configuration, configuration wizard needs to be relaunched. The steps involved are as follows:

- a. Select **Maintenance > System**.
- b. Select **Configuration System**.

Best practices

AutoSupport (ASUP) configuration

Enabling AutoSupport is recommended in DDVE. Although Secure Remote Services is not yet supported in AWS, an email transfer server can be used to transfer ASUP files.

AWS licensing

The DDVE license is node locked, which means the same license cannot be used on multiple DDVE instances. To facilitate DDVE license management, it is recommended to use served-mode licenses if multiple DDVEs are to be deployed.

Storage best practices

Use the appropriate storage type

Use GP2 EBS volumes for the root disk, NVRAM disk, and metadata disks.

Object storage specifications

Metadata disk storage is recommended to be 10% of the total capacity. Each metadata disk is recommended to be 1 TiB.

Block storage specifications

DDVE with block storage supports a maximum capacity of 16 TB. The recommended size of each data disk is 1 TiB.

Metadata disk storage expansion

The metadata storage recommendation in Object storage specifications is based on 10X deduplication ratio and 2X compression. For workloads with a higher deduplication ratio, more metadata storage may be required. If metadata storage usage exceeds 80%, an alert is generated. Add a metadata disk to the DDVE immediately to avoid running out of space.

Spindle group

It is not required to specify a spindle group when adding metadata disks. The spindle group assignment is balanced automatically when adding storage. Do not set or change the spindle group settings manually. Run the `storage show all` command to verify that each data volume is assigned to a different spindle group.

Object storage bucket configuration

- The bucket that is provided during file system creation must be empty, otherwise file system creation fails.
- When the file system is destroyed, the associated bucket and the objects it contains are not automatically deleted or removed. The bucket must be intentionally deleted to avoid incurring the cost for the content stored in the bucket.
- Do not configure any life-cycle policy on the bucket as it might result in loss of critical data.

Converting from evaluation to production

Rather than converting an evaluation version of DDVE to a production version, Dell Technologies recommends a fresh deployment. If it is required to convert from an evaluation version to production version, Dell Technologies recommends the following:

- Destroy the existing file system
- Delete any small data disk (not the root or NVRAM disks)
- Configure new disks according to the recommendations in this guide

Operational best practices

AWS recommends that AWS CloudTrail logs must be enabled to enable governance, compliance, and operational and risk auditing of the AWS account. AWS CloudTrail allows the following:

- View the event history of AWS cloud activity, including AWS Management Console actions, AWS SDKs, CLI, and other AWS services.
- Identify the initiator of actions, resources involved, and event timing.

This event history helps to simplify security analysis, resource change tracking, and troubleshooting.

Conclusion

Security best practices

- Avoid public IP address to configure DDVE.
- For secure access, it is recommended to disable the username/password based user authentication. If the username/password based authentication is wanted, it is better to use a stronger password.
- Since the DDVE in AWS is always running in a VPC, the VPC should be configured so that only required and trusted clients have access to the DD system.

Network best practices

- It is recommended to use public or private subnet architecture to deploy the DDVE in private subnet.
- It is highly recommended to use VPN connections between different geographical regions (VPCs).
- The DDVE object store feature needs connectivity to its object storage, such as to the S3 bucket. The object store communication is over https, so the outbound security group setting must allow communication over port 443. There are different ways to enable DDVE connectivity to the object store and the recommended one is using the VPC endpoint for accessing the Amazon S3.

Conclusion

DDVE can be easily deployed on AWS platform and can protect the applications running on the cloud environments. DDVE can be up and running in minutes and delivers increased transactional and operational efficiencies along with high-speed and variable length deduplication.

Technical support and resources

Dell.com/support is focused on meeting customer needs with proven services and support.

Related resources include:

- [DDVE on AWS Installation and Administration Guide](#)
- [DDVE Installation and Administration Guide](#)
- [DDVE in AWS and VMC Best Practices Guide](#)

Dell PowerProtect Data Manager: File System Backup and Recovery

Abstract

This white paper focuses on file system backup and recovery using Dell PowerProtect Data Manager.

October 2022

Revisions

Date	Description
July 2019	Initial release
September 2019	Document revised for PowerProtect Data Manager version 19.2 release
December 2019	Document revised for PowerProtect Data Manager version 19.3 release
February 2021	Document revised for PowerProtect Data Manager version 19.7 release
May 2021	Document revised for PowerProtect Data Manager version 19.8 release
October 2021	Document revised for PowerProtect Data Manager version 19.9 release
May 2022	Document revised for PowerProtect Data Manager version 19.10 release
October 2022	Document revised for PowerProtect Data Manager version 19.12 release

Acknowledgments

Author: Vinod Kumar Kumaresan

The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

This document may contain certain words that are not consistent with Dell's current language guidelines. Dell plans to update the document over subsequent future releases to revise these words accordingly.

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

Copyright © 2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [10/27/2022] [Technical White Paper] [H18659.7]

Table of contents

Revisions	2
Acknowledgments	2
Table of contents	3
Executive summary	4
Audience	4
1 Introduction to Data Manager for File System	5
1.1 Data Manager key features for file system protection and recovery	5
2 File System backup technology and models in Data Manager	6
2.1 Block-based backup technology	6
2.2 File-based backup technology	7
2.3 Data Manager models for file system protection and recovery	7
2.4 Roadmap to protect a file system	9
3 File System backup configurations	10
3.1 File System Agent	10
3.2 Enabling File System asset source	11
3.3 File System host configuration	11
3.4 File system asset discovery	12
3.5 Protecting Windows clustered disks with Data Manager	13
4 File system protection policy	15
4.1 Protection rule for file system	16
4.2 Exclusion filters for file system data protection	16
4.3 File system parallel backup settings	17
5 Data Manager File System Backup	18
5.1 Centralized file system backup workflow	18
5.2 Self-service file system backup workflow	19
6 Data Manager File System Restore	20
6.1 Centralized file system restore workflow	20
6.2 Self-service file system restore workflow	21
6.3 Centralized file-level restore of file systems	22
6.4 Search support for file system workloads	22
6.5 Self-service file-level restore of file systems	24
A Technical support and resources	25
A.1 Related resources	25

Executive summary

Business Case: Challenges

Today's data protection is either too complex, requires multiple vendors, does not scale, or does not meet the needs of fast-growing, modern, and agile organizations of all sizes. As businesses continue to consume IT resources differently, there is a need for powerful, efficient, and trusted data protection to enable organizations to transform to meet future demands when modernizing their IT environment. One of the biggest challenges is determining how we can turn the data we protect and manage into value. Customers are challenged with backups, recovery, and ultimately the management and governance of the data they are protecting.

Solution Overview: Why PowerProtect Data Manager?

Dell PowerProtect Data Manager is the next generation data management platform to transform traditional data protection to comprehensive data management. Data Manager gives IT the trusted data protection they know from Dell Technologies, combined with operational simplicity that protects workloads and file systems running on-premises with self-service capabilities for operational efficiency and IT governance controls to ensure compliance. SaaS-based management makes it easy to monitor, analyze, and troubleshoot distributed data protection environments from anywhere.



Data Manager gives valuable insight into protected on-premises and in-cloud workloads, applications, file systems, and virtual machines. Designed with operational simplicity and agility in mind, Data Manager enables the protection of traditional workloads including Oracle, Exchange, SQL, SAP HANA, NAS, file systems, Kubernetes containers, and virtual environments.

This white paper focuses on file system backup and recovery with Data Manager, which ensures reliable and efficient data protection functionality. It also describes the file system backup architecture, backup and recovery workflows, and deployment requirements.

Audience

This white paper is intended for Dell Technologies customers, partners, and employees who are looking for file system data protection and management using Data Manager.

1 Introduction to Data Manager for File System

Data Manager offers centralized oversight of all protected file system copies. This makes it simple to track and enforce service level objective (SLO) compliance for backup and recovery, RPOs, and Storage retention lock. Data Manager discovers copies sent to protection storage, then catalogs and makes protection copies available for compliance measurement to ensure protection compliance and quality of service.

The File System Agent allows an application administrator to protect and recover data on the file system host. Data Manager integrates with the File System Agent to check and monitor backup compliance against protection policies. Data Manager File System Agent has been designed to support the file system backup, restore, and replication workflows. This white paper describes how effectively you can protect file systems using Data Manager with Dell PowerProtect DD series appliances as target storage.

1.1 Data Manager key features for file system protection and recovery

- Centralized file system backup and recovery (volume and file level) through Data Manager
- Self-service file system backup and file-level recovery (volume and file level) using command-line utility
- Block-based and File-based file system support for file system workloads. For more details on file system workloads compatibility, see section "File Systems" in the [PowerProtect Data Manager support matrix](#)
- Supports centralized file level restores of block-based file system backups.
- File level restore from Data Manager UI
 - Support restore to original or alternate host for file-based backups
 - Creation of separate directory to separate the restore files
- Support for backup of Non-LVM or physical disks
- Supports exclusion of files and folders from file system backups through use of file exclusion
- Data Manager provides support to run file system backups in parallel to reduce the time taken for backups
- No excessive metadata generation and PowerProtect DD series appliances Active Tier storage consumption has been optimized for file-based backups
- Automated host agent configuration during policy creation
- Flexibility to define and assign exclusion filter to file system protection policy for excluding certain files and folders
- **Encryption over the wire (EOW)** of file system backup and restore data - Data Manager provides EOW of data during some backup and restore operations. If enabled, encryption is automatically applied to file system workloads
- Data Manager File System Agent supports the Microsoft System State Recovery (SSR), Active Directory Restore, and bare-metal recovery (BMR) disaster-recovery mechanisms
- Data Manager File System agent now allows the selection of clustered drives and logical cluster hosts as assets and asset sources
- Starting with PowerProtect Data Manager 19.12, file indexing and search is supported for file system workloads

2 File System backup technology and models in Data Manager

2.1 Block-based backup technology

Data Manager performs a file system level full and incremental backup using block-based backup (BBB) technology. During the backup, the application agent scans a volume or a disk in a file system and backs up all the blocks that are in use in the file system. Unlike the traditional file system backup, BBB supports high-performance backups with a predictable backup window.

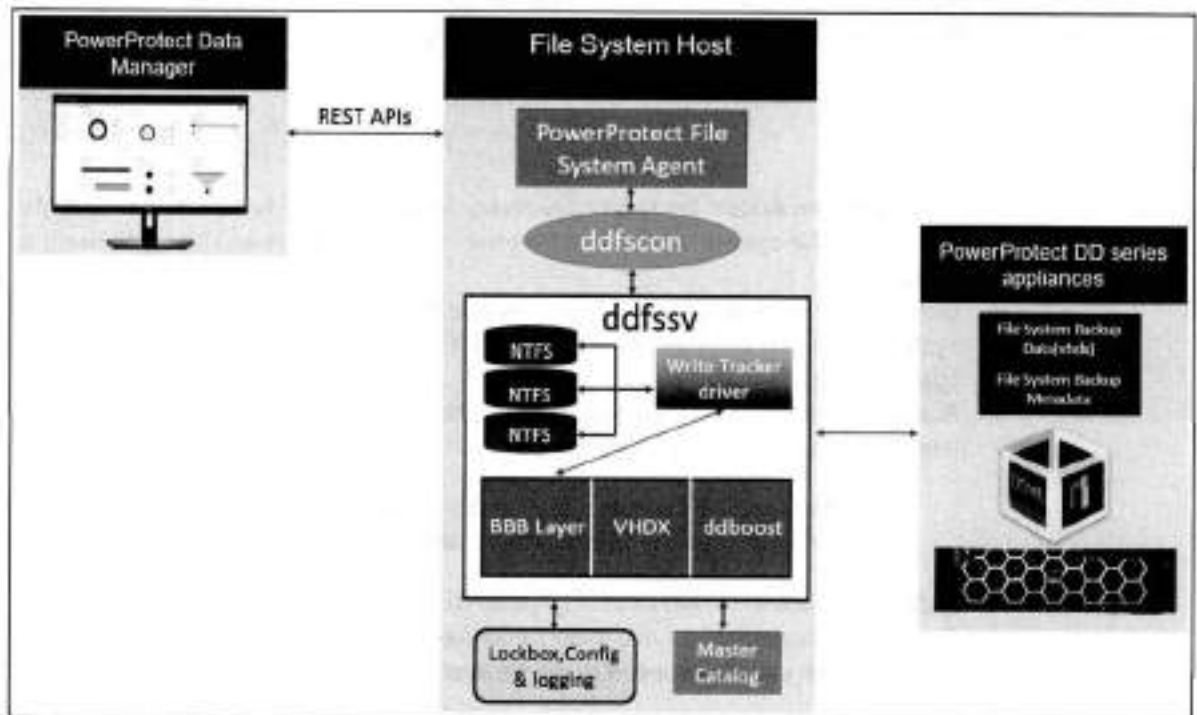


Figure 1: Data Manager Block-Based Backup Technology:

Block-based backups provide instant access to the backups. The block-based backups enable you to mount the backups by using the same file systems that you used to backup the data.

The File System Agent's block-based backups support the following capabilities:

- Mounting of a backup as a file system
- Mounting of an incremental backup
- Sparse backup support
- Backups of operating system-deduplicated file systems as source volumes on Windows
- Forever virtual full backups to DD series appliances
- DD retention lock
- Recoveries from DD series appliances

Block-based backups are useful for datasets that are under 10 TB with a single volume under 5 TB, and a daily change rate under 5%.

2.2 File-based backup technology

File-based backups (FBB) traverse through the entire directory structure of the file system to backup all the files in each directory of the file system. FBBs can provide additional capabilities such as exclusion. These backups take longer to complete when compared to block-based backups. The File System Agent performs a FBB of the protected assets when an exclusion filter is applied to a protection policy.

Note: Exclusion filters cannot be applied to self-service protection policies and to backups taken through self-service CLI.

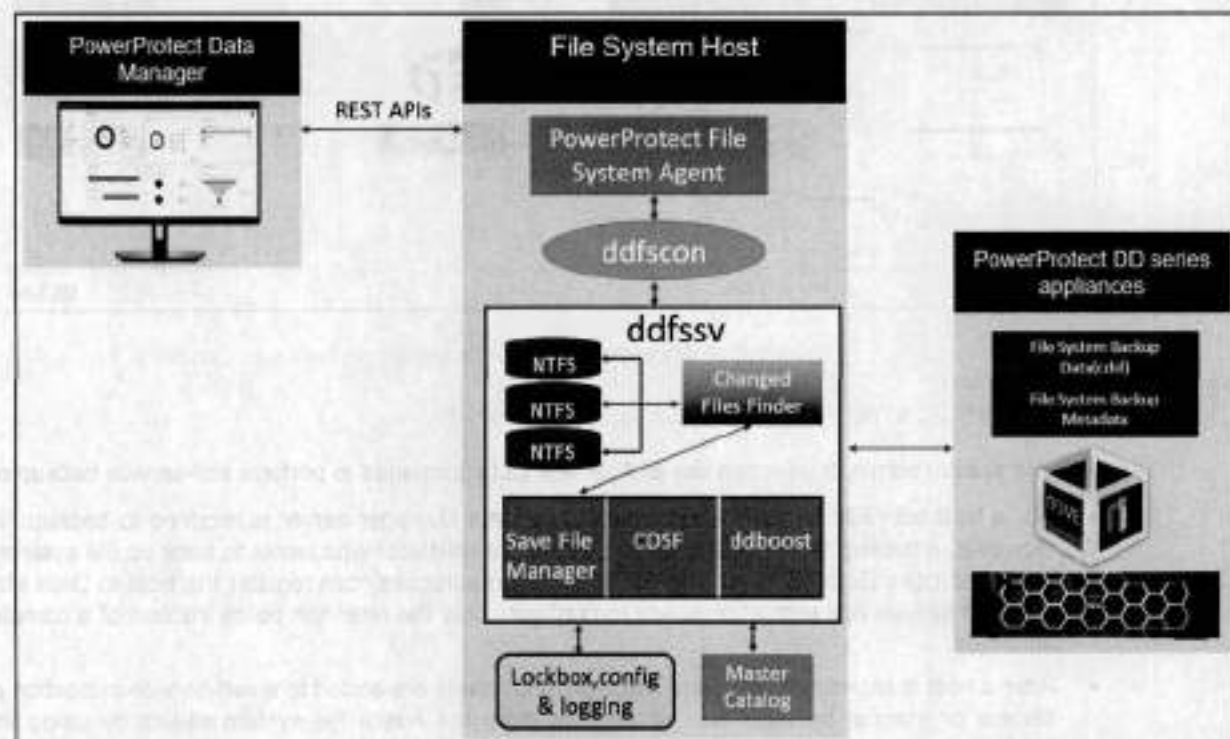


Figure 2: Data Manager File-Based Backup Technology

2.3 Data Manager models for file system protection and recovery

File system protection data zone components include Data Manager server, PowerProtect File System Agents, file server host, and storage. Data Manager for file system protection has been built with two service model features:

- Self-service file system protection and recovery
- Centralized file system protection and recovery

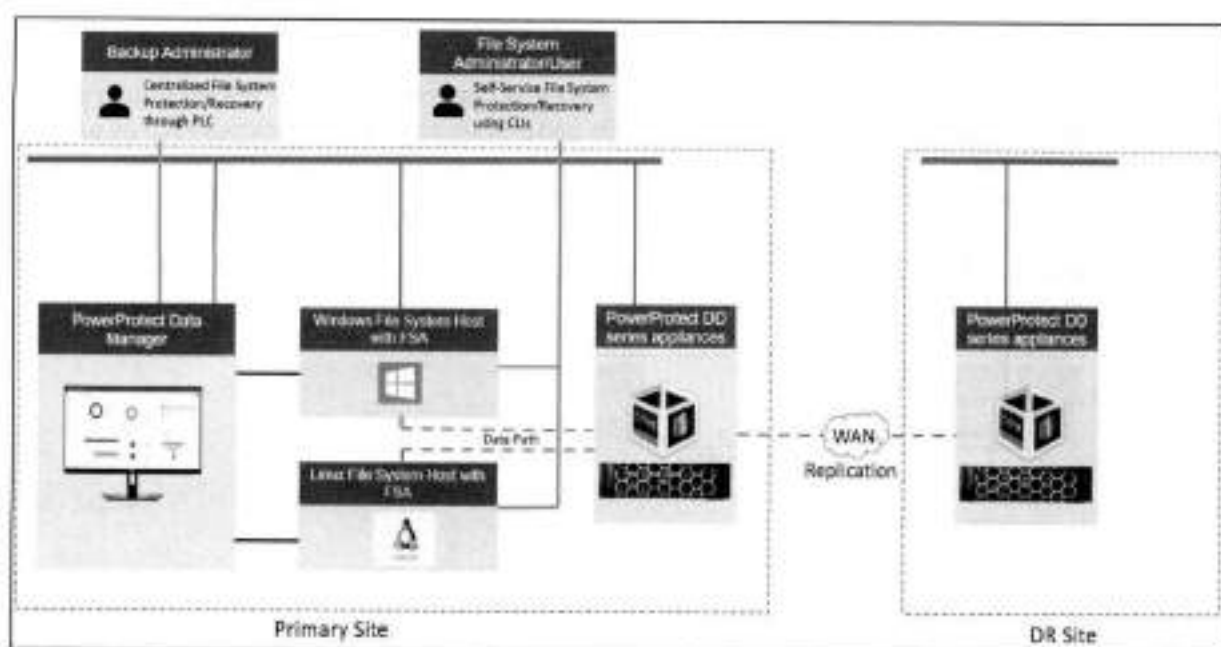


Figure 3: Data Manager models for file system protection and recovery

a) Self-service file system protection and recovery

- File system admin or user can use self-service CLI commands to perform self-service backup or restore
- For a host with File System Agent installed, a Data Manager server is required to backup file systems. However, a backup administrator or file system administrator who wants to back up file systems manually (and who uses Data Manager only for compliance purposes) can register the host to Data Manager and create a self-service protection policy to configure only the retention policy instead of a complete backup schedule.
- After a host is registered with Data Manager and assets are added to a self-service protection policy, self-service or manual backups can be performed on the host's file system assets by using the (ddfssv) command-line utility.
- Self-service restore can restore from backups that were centralized or self-service and can be done to a local or remote server.

Note: To enable self-service protection, self-service protection option is selected when creating the file system protection policy in the Data Manager.

b) Centralized file system protection and recovery

- Centralized file system model is built for backup administrators to perform policy-based file system backup, recovery, replication, and long-term retention of copies. With the centralized protection feature, Data Manager manages the entire file system backup workflow, including the schedule.
- Centralized protection is supported through protection policy. After the File System Agent is installed on the client, the client is auto discovered on Data Manager and enables the administrator to approve the client.
- Choosing the centralized protection option during protection policy creation enables Data Manager to manage all protection centrally.

2.4 Roadmap to protect a file system

The following roadmap provides the steps required to configure the File System Agent in Data Manager to run protection policies.

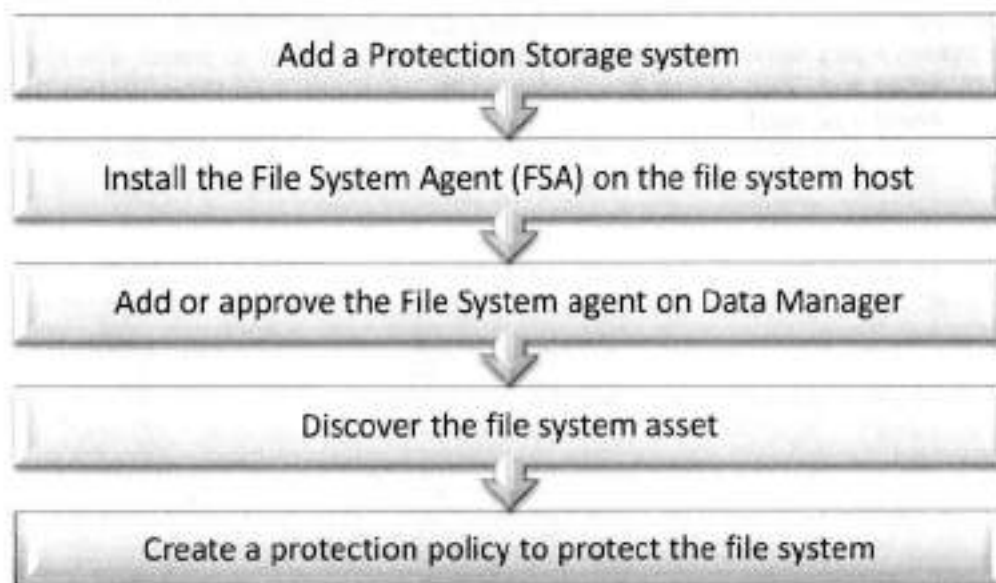


Figure 4: Roadmap to protect a file system

See [PowerProtect Data Manager File System Agent User Guide](#) for details about how to install, configure, and perform a file system backup/restore.

3 File System backup configurations

3.1 File System Agent

The File System Agent needs to be installed on the host that is planned to protect. The File System Agent binaries (Windows and Linux) can be downloaded on the below path from **PowerProtect Data Manager** → **Settings** → **Agent downloads**.



See *PowerProtect Data Manager File System Agent User Guide* for details on how to install File System Agent on supported Windows and Linux hosts.

Data Manager supports the coexistence of agents on the same Windows or Linux host for the following:

- Microsoft SQL agent and the File System Agent on Windows
- Microsoft Exchange agent and the File System Agent on Windows
- Oracle RMAN agent and the File System Agent on Linux
- SAP HANA agent and the File System Agent on Linux

Software compatibility information for the Data Manager software and the File System Agent is provided in the eLab navigator, available at <https://elabnavigator.dell.com/eln/modernHomeDataProtection>.

3.2 Enabling File System asset source

By enabling the file system asset source, Data Manager can protect and recover file system data when Data Manager is integrated with the File System Agent.



3.3 File System host configuration

After the File System Agent is installed on the file system host, File System Agent can be added, approved, and rejected for the pending agent requests. Select **Infrastructure > Application Agents** and click **Add** or **Approve** as required.

The "Auto Allow List" all option is disabled by default. When enabled, all preapproved application agents are automatically approved.

File System Agent registered with Data Manager:



File System backup configurations

On successful registration, the file system host is listed on the **Infrastructure > Asset Sources** section. Asset discovery is initiated by default after registration of the file system host to Data Manager.

File system host is discovered in the **Asset Sources > File System** section.



3.4 File system asset discovery

Discovered file system assets in the **Infrastructure > Assets** section.



3.5 Protecting Windows clustered disks with Data Manager

Starting with Data Manager version 19.10, the File System Agent provides support for protecting the Windows clustered file system. With this feature, customers can use the File System Agent to protect their clustered disks, like the regular file system drives.

Both BBB and FBB are supported for Windows clustered disk protection. During the failover, the backup will continue through the node that owns the cluster disks. This prevents administrators from having to manually reconfigure for data protection continuity.

The following figure shows the **Application Agents** view for the Windows cluster nodes registered with Data Manager.



The above cluster nodes registered with Data Manager discover the clustered disks as assets, as shown below.



File System backup configurations

A file system protection policy can be created to back up the cluster assets, in the same way of protecting the regular file system drives.



Backup copies available for Windows cluster assets.



Once the backup is successful, Data Manager provides the option to perform FLR and Image level restore to the original, or to an alternate location, from the backup copy.



4 File system protection policy

A file system protection policy can be created from Data Manager UI to protect file system data.

Protection life cycle policy defines a set of objectives that apply for a specific duration. Data Manager provides centralized and self-service protection options to specify one of the following "Purposes" for the protection policy to back up Linux/Windows file systems. These objectives drive configuration, active protection, and data management operations that satisfy Service Level Agreements (SLAs).

For file system protection, you can select one of three types:



- **Centralized Protection** - To use Data Manager to manage all protection centrally.
- **Self-Service Protection** - To use the file system to create local backup protection. Data Manager creates a protection policy and manages extra stages.
- **Exclusion** - If there are assets within the protection policy to exclude from data protection operations.

Option to enable indexing for file search and restore from the protection policy, when one or more search nodes are deployed with PowerProtect Data Manager.



File system protection policy

After the policy configuration is complete, an informational message appears to confirm that Data Manager has saved the protection policy. Initially there will be two configuration tasks running:

- The first task will create a storage-unit on DD series appliance.
- The second task will update the agent lockbox on the file system host and add the new storage-unit credentials.

Job ID	Status	Description	Job Type
C0KTYL4RH	Success	Configuring File System - Filesystem_Backup	Config
8DVM2YSD	Success	Performing Policy Configuration - Filesystem_Backup - PROTECTION	Config

See [PowerProtect Data Manager File System Agent User Guide](#) for detailed steps on how to create a protection policy for file system backup.

4.1 Protection rule for file system

A protection rule automatically determines which assets get assigned to protection policies when the assets are discovered. A protection policy must exist prior to creating the dynamic filter. An asset can only belong to one protection policy.



4.2 Exclusion filters for file system data protection

Data Manager provides the option to exclude data (file system files and folders) from assets that are assigned to protection policies. File system exclusion filters can be defined and applied to a protection policy to exclude certain files and folders from file system data protection.

An exclusion filter provides the following options for excluding file and folders from a file system backup.

The screenshot shows the 'New Information' form in the PowerProtect Data Manager interface. The form is titled 'New Information' and contains several input fields and sections:

- Name:** A text input field with the value 'FileSystem'.
- Source:** A text input field.
- Filter Conditions:** A section with four filter types:
 - File Type:** A dropdown menu set to 'Application' and a text input field with '*.txt,*.xlsx,*.pdf'.
 - Modified Time:** A text input field with '1/1/2024'.
 - File Size:** A text input field with '100'.
 - Folder Path:** A text input field with 'C:\Program Files\Microsoft Office\Office16'.
- Exclusion Filter:** A section with a text input field containing 'Microsoft Word'.
- Buttons:** An 'Add Filter' button is located at the bottom right of the form.

Exclusion Filter Conditions	
Conditions	Description
File Type	For example - .txt, .xlsx, .pdf
Modified Time	File/Folder modification time
File Size	File/Folder size
Folder Path	File/Folder path

After the filters are created, filters can be applied during the new file system protection policy creation or it can be applied to an existing file system protection policy. For verification, backup logs provide details of the files and folders which are excluded from backup according to the filter defined.

4.3 File system parallel backup settings

Data Manager enables running file system backups in parallel to reduce the time taken for backups. This setting defines the maximum concurrent network sessions from the client to DD series appliance at any given time. The number of streams to use for the backup can be specified in the configuration file `.ddfssv.fsagentconfig` or through the self-service CLI. However, it is best to set the parallelism value in the configuration file because the parallelism value provided in the configuration file takes precedence over the parallelism value that is provided in the CLI.

Note: Backup parallelism is only available on supported Windows systems. Because the parallelism setting is defined at the host level, the parallelism setting must be set on every Windows host where parallel file system backups are enabled. This value must be an integer. The default value is 8.

See [PowerProtect Data Manager File System Agent User Guide](#) for details on how to specify the number of streams to use for the backup in the `.ddfssv.fsagentconfig` file in the `C:\Program Files\DPSAPPS\fsagent\settings` directory on the file system host or by using the command-line option.

5 Data Manager File System Backup

Data Manager enables discovering, managing, monitoring data protection, and replicating file system assets through integration with the File System Agent. File system assets are protected in Data Manager with centralized and self-service file system protection features.

Data Manager self-service protection enables users to perform backup and restore using a self-service CLI workflow for Windows and Linux assets. With agility and self-service feature, data owners can perform backup and recovery within native applications.

5.1 Centralized file system backup workflow

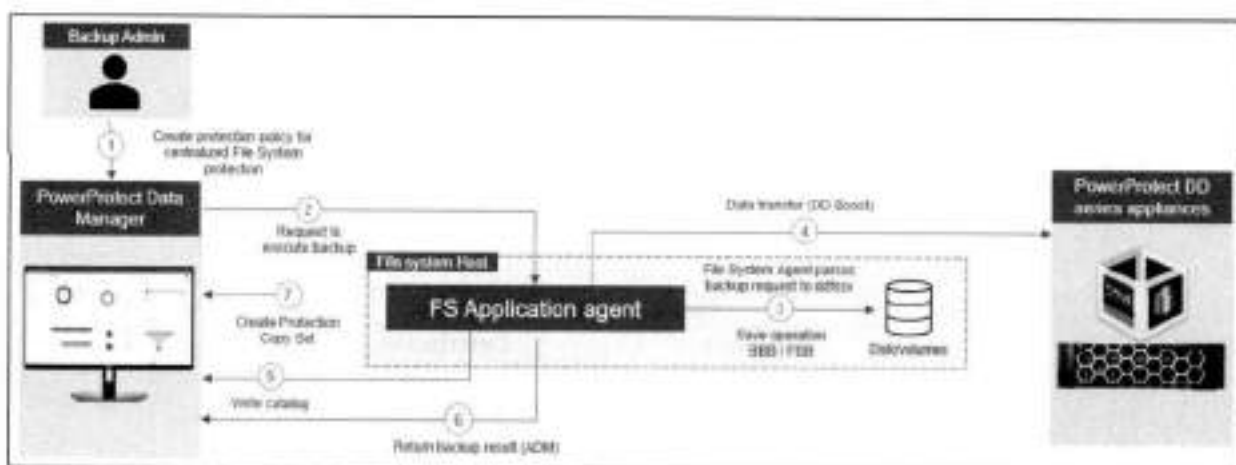


Figure 5: Centralized file system backup workflow

1. Protection policy is created for centralized File System protection.
2. At the time of scheduled backup, the Data Manager requests the File System Agent to perform save operation for the file system data.
3. File System Agent parses the backup job request and converts into (ddfsav) utility commands to perform save operation.
4. File System Agent verifies DD series appliance connectivity and writes the file system data to the storage-unit created on the DD series appliance. (Waits for 15 minutes)
5. File System Agent writes the catalog details to catalog database on Data Manager.
6. ADM agent retrieves the result from FS agent and updates the backup status to Data Manager.
7. Creates and maintains Protection Copy Set (PCS) in Elasticsearch database

5.2 Self-service file system backup workflow

A host with the File System Agent installed requires Data Manager to back up the file systems. To back up file systems manually and use Data Manager, the file system host needs to be registered with Data Manager and a self-service protection policy needs to be created. Data Manager discovers these backups and enables centralized restore operations. You can also perform a manual restore operation.

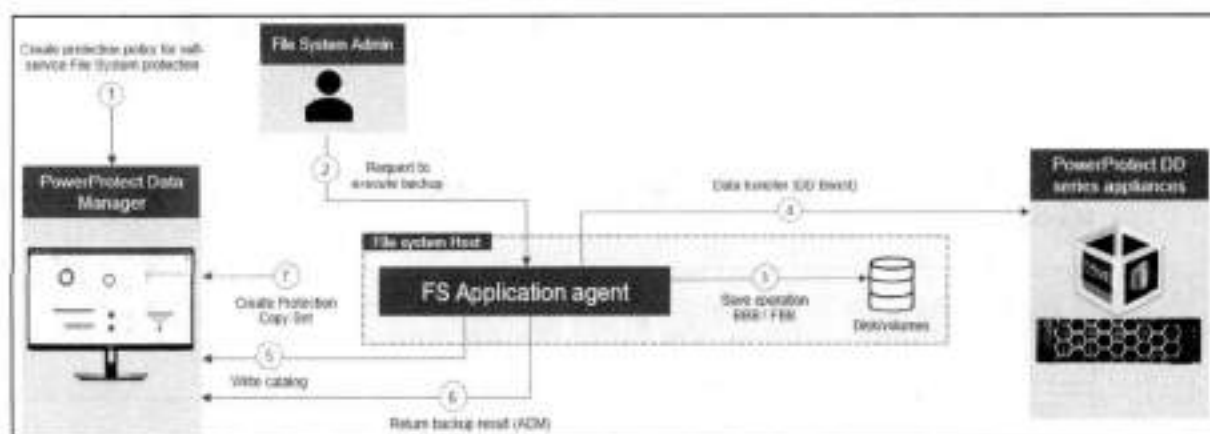


Figure 6: Self-service file system backup workflow

1. Protection policy is created for self-service File System protection.
2. File system administrator launches the (ddfsv) utility using command line on the file system host and inputs the below details to initiate backup.
Backup schedule (Full, Incremental)
Storage IP address
Storage Username
Storage password
3. (ddfsv) utility performs save operation.
4. File System Agent verifies DD series appliance connectivity and writes the file system data to the storage-unit created on the DD series appliance.
(Waits for 15 minutes)
5. File System Agent writes the catalog details to catalog database on Data Manager.
6. ADM agent retrieves the result from FS agent and updates the backup status to Data Manager.
7. Creates and maintains Protection Copy Set (PCS) in Elasticsearch database

6 Data Manager File System Restore

When file systems are protected within a protection policy in a Data Manager, the file system data can be recovered using the centralized restore functionality or by directly using the self-service restore feature.

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 enough space available on the target volume for the restore

6.1 Centralized file system restore workflow

A file system host image-level restore allows recovering data from backups of file systems performed in Data Manager.

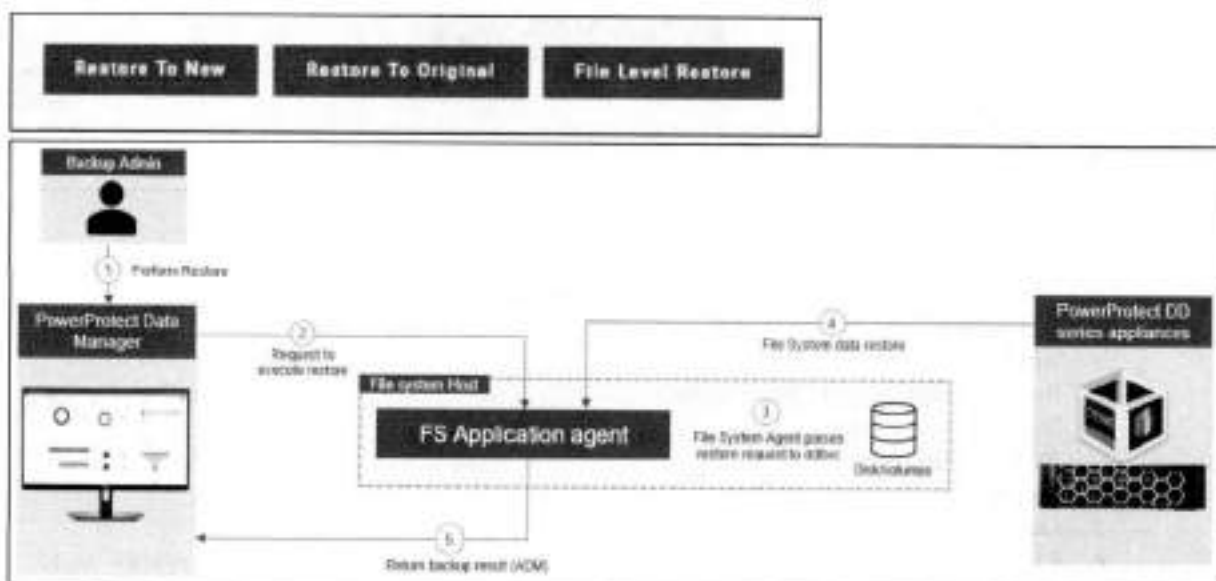


Figure 7: Centralized file system restore workflow

1. Backup administrator creates the recovery job on Data Manager UI with below UI inputs,
 - Source file system backup
 - Destination file system
 - Restore options
 - Restore file location
2. Data Manager requests its ADM agent to dispatch the restore operation to File System Agent.
3. File System Agent parses the recovery job request and converts into (ddfsrc) utility commands and executes the restore operation based on inputs provided.
4. File System Agent verifies DD series appliance connectivity and the requested file system data is restored to the specified destination.
5. ADM agent retrieves the result from File System Agent and updates the restore status to Data Manager.

Note: If the destination file system asset already contains some data, this data will be overwritten.

6.2 Self-service file system restore workflow

Self-service image-level restores of file systems can be performed by using the `ddfsrc` 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

Before running the (`ddfsrc`) command to perform a self-service image-level restore of file Systems, the (`ddfsadmin`) backup command can be used to query a list of all the local backups taken for a host and obtain the ID of the save set for restore. Specify the ID of the save set as an input to the `ddfsrc` command. If restoring to the original host, the password will be picked up from the lockbox.

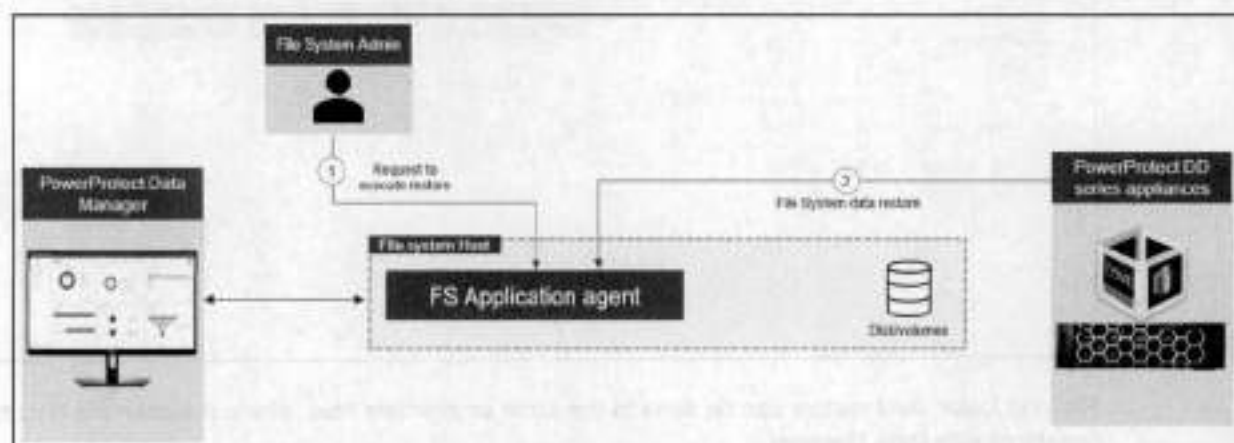


Figure 8: Self-service file system restore workflow

1. File system administrator launches the `ddfsrc` command-line utility and below details to be entered as input. Once the inputs are validated and executed the `ddfsrc` utility performs recover operation.
 - Source file system backup
 - Destination file system
 - Restore options
 - Restore file Location
2. File System Agent verifies DD series appliance connectivity and the requested file system data is restored to the specified destination.

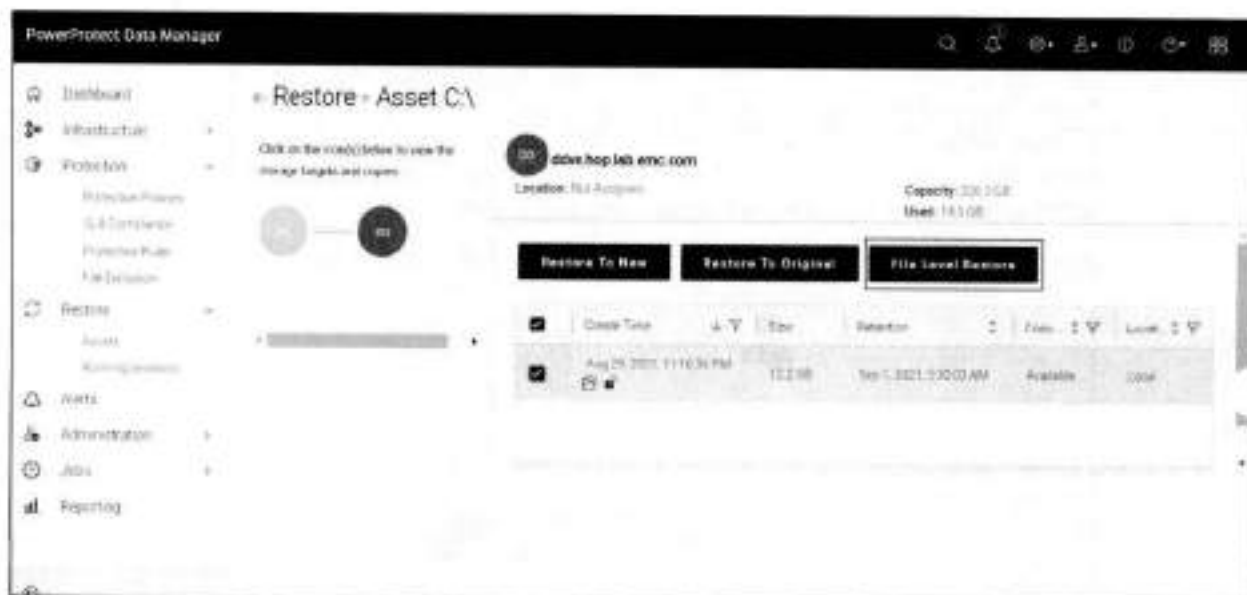
The self-service restore feature provides the following options to restore the data:

- Restore to same host and same location
- Restore to same host and different location
- Restore to different host and location

6.3 Centralized file-level restore of file systems

Data Manager provides the option to restore files and folders from file system backup through Data Manager centralized console.

Recovery > Assets > File Level Restore



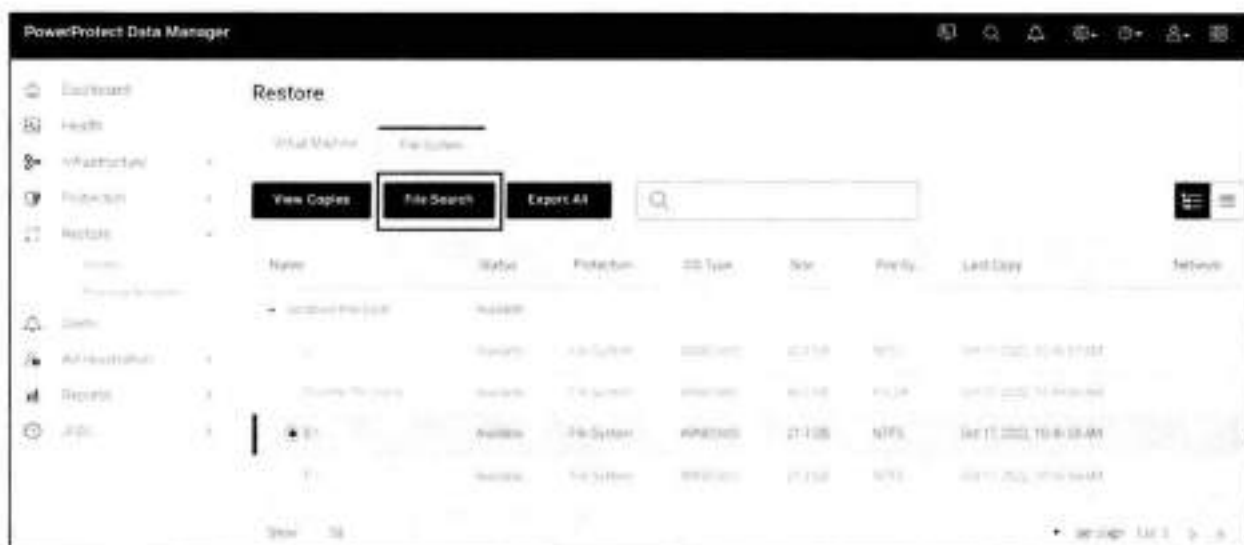
File and folder level restore can be done to the same or alternate host, where the alternate host needs to be registered with Data Manager.



6.4 Search support for file system workloads

Starting with PowerProtect Data Manager 19.12, file indexing and search is supported for file system workloads. As a prerequisite, one or more search nodes must be deployed. Indexing for file search and the restore option is enabled from the protection policy.

When the file system backup is completed, the **File Search** option is available to search and restore the required files.

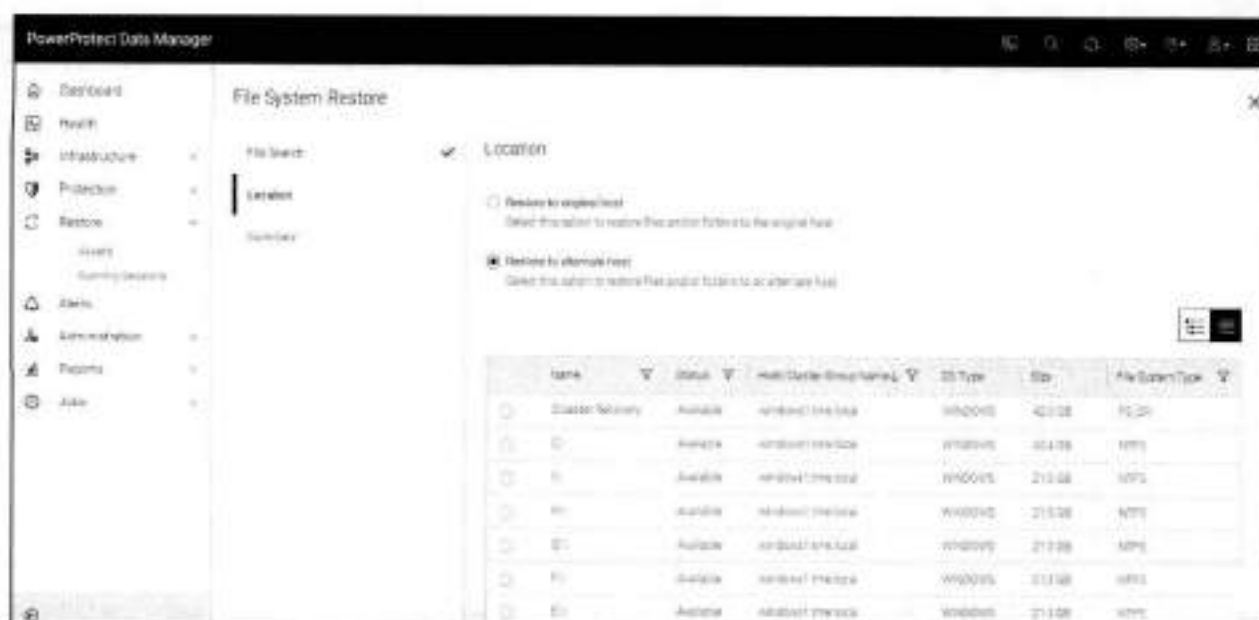


The "FS File Search" window provides the ability to search for the required files using the different search criteria options.



File search and restore options:

- Restore to original host - Select this option to restore files and/or folders to the original host.
- Restore to alternate host - Select this option to restore files and/or folders to an alternate host.



6.5 Self-service file-level restore of file systems

Self-service file-level restores of file systems can be performed using the `ddfsrc` command with the `-l` option. A file that contains the list of file(s) to be restored is created before executing the command. The location of this file is given as an input to the `-l` option.

For more details on performing self-service file-level restore of file systems, see [PowerProtect Data Manager File System Agent User Guide](#).

A Technical support and resources

Dell.com/support is focused on meeting customer needs with proven services and support.

The [Data Protection Info Hub](#) provides expertise that helps to ensure customer success with Dell data protection products.

A.1 Related resources

- [PowerProtect Data Manager File System User Guide](#)
- [PowerProtect Data Manager Administration and User Guide](#)
- [PowerProtect Data Manager Deployment Guide](#)
- [PowerProtect Data Manager Compatibility Matrix](#)

Dell PowerProtect Data Manager: Virtual Machine Backup and Recovery

July 2023

H18660.8

White Paper

Abstract

This white paper focuses on using Dell PowerProtect Data Manager to protect virtual machine environments.

Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2019–2023 Dell Inc. or its subsidiaries. Published in the USA July 2023 H18660.8.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Contents

Executive summary	4
Introduction	6
Deployment requirements	13
Data Manager protection policy options for virtual machine protection	14
PowerProtect Search Engine	18
Virtual machine consistent backup workflows	25
Data Manager use cases for virtual machine recovery	28
Disaster recovery	35
Technical support and resources	37

Executive summary

Overview

Dell PowerProtect Data Manager provides software-defined data protection, automated discovery, deduplication, operational agility, self-service, and IT governance for physical, virtual, and cloud environments. PowerProtect Data Manager enables users to:

- Orchestrate protection directly through an intuitive interface or empower data owners to perform self-service backup and restore operations from their native applications
- Ensure compliance and meet even the strictest of service level objectives
- Leverage existing Dell PowerProtect appliances

PowerProtect Data Manager provides a consistent protection experience for VMware environments. Data Manager is the only solution to provide native vSphere integration with vCenter for virtual machine protection, offering storage and backup admins, and virtual machine owners to choose a storage policy to apply to every virtual machine automatically when it is instantiated.

Audience

This white paper is intended for Dell customers, partners, and employees looking to understand how Data Manager helps to protect VMware workloads.

Revisions

Note: This white paper is based on Data Manager version 19.13 release. The contents of this white paper are updated for each version release.

Date	Part number/ revision	Description
July 2019		Initial release
December 2019		Update for PowerProtect Data Manager version 19.3
February 2021		Update for PowerProtect Data Manager version 19.7
May 2021		Update for PowerProtect Data Manager version 19.8
September 2021		Update for PowerProtect Data Manager version 19.9
May 2022		Update for PowerProtect Data Manager version 19.10
July 2022		Update for PowerProtect Data Manager version 19.11
October 2022		Update for PowerProtect Data Manager version 19.12
April 2023	H18660.6	Update for PowerProtect Data Manager version 19.13
April 2023	H18660.7	Editorial (template) update only
July 2023	H18660.8	Update for PowerProtect Data Manager version 19.14

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

Author: Chetan Padhy

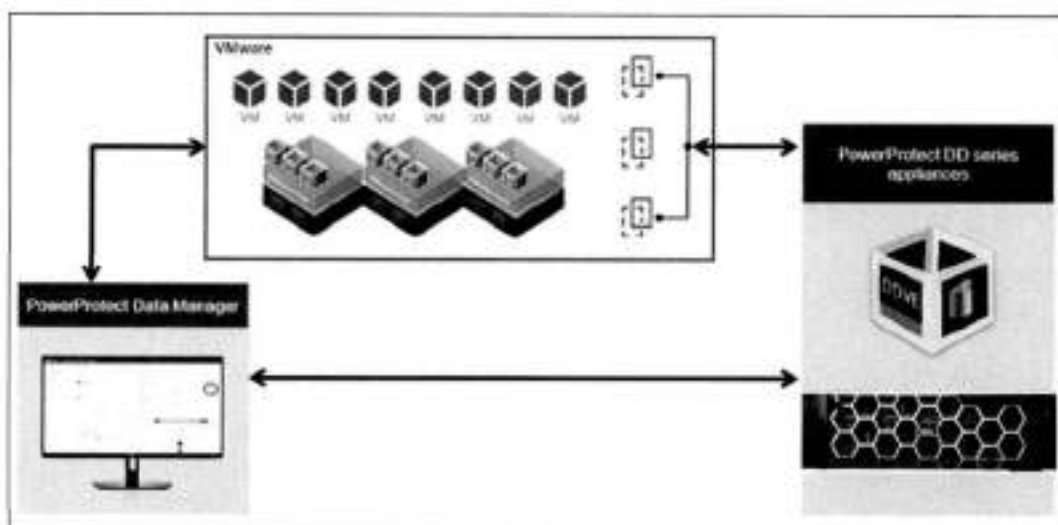
Note: For links to other documentation for this topic, see the [Data Protection Info Hub](#).

Introduction

Data Manager integration with virtual machines

Data Manager integration with virtual machines helps to manage, protect, and reuse virtual machine data across the enterprise by deploying services to accomplish the following tasks:

- Discover, access, and recover virtual machine copies nondisruptively across primary and protection storage without introducing new infrastructure or complexity
- Automate efficient copy creation
- Efficiently automate data retention SLA compliance, ensuring that the right number of copies are stored in the right place at the right level of protection
- Optimize operations based on actionable analytics and insight



Data Manager features for virtual machine backup and recovery

Data Manager includes the following features for virtual machine backup and recovery:

- Supports governance for VMware virtual machines using VM Direct Engine with advanced restore capabilities, such as instant access and instant restore. VM Direct Engine enables you to:
 - Monitor and manage adherence to retention policies for backups
 - Create ad hoc backups of VMware virtual machines
 - Restore VMware virtual machine copies, including to original, to alternate, instant access, instant restore, file level restore, and restore to ESXi host
- Supports vSphere Client plug-in for PowerProtect, which provides a subset of Data Manager functionality within the vSphere Client, including restore to the original virtual machine, restore to a new virtual machine, and instant access virtual machine restore
- Supports virtual machine disk-level (VMDK) restore to the original location
- Supports VMware high-availability deployment configuration
- Supports VM file indexing and restore from search results

- Supports VM disk exclusion for backup and restore
- Supports on-demand backups of virtual machines in the vSphere Client plug-in
- Provides data protection for VMware Cloud virtual machines
- Includes Data Manager Search software. Rapidly search on specific file attributes across indexed virtual machine backups
- Perform file-level restores from Data Manager virtual machine backup copies by using the PowerProtect portlet in the vSphere Client
- Support for multiple backup optimization modes Capacity to optimize for storage consumption, and Performance to optimize for speed
- An option to enable warning and failure thresholds for backup and restore operations when the available space on the datastore is low
- Option to improve backup performance by excluding memory and swap files (C:\swapfile.sys, C:\pagefile.sys, and C:\hiberfil.sys) from Microsoft Windows virtual machine backups
- The option to restore the tags and categories associated with virtual machines when restoring a virtual machine protection policy backup
- For virtual machine protection jobs with a daily, weekly, or monthly schedule, a failed job or task can be restarted in Data Manager by configuring auto retry in the entrypoint.sh file
- VMware Storage Policy Based Management (SPBM) policies can be paired with Data Manager protection policies, allowing to manage all virtual machine storage and protection policies from the vSphere Client
- The quiescing of virtual machines can be enabled or disabled with APIs.
- Option to delete a search node without data loss using APIs
- The deletion of a search node from multinode search clusters from Data Manager UI
- Enable users to select multiple files across multiple copies and virtual machines with VM search while performing File-level restore
- DD Boost compressed restore—reduction in restore times for bandwidth limited environments
- Transparent Snapshots Data Mover (TSDM) protection mechanism for performing crash-consistent and SQL application-aware virtual machine protection
- Option to change the virtual machine network settings during restore.
- Override User Account Control (UAC) option is added for Windows and Linux file-level restores (FLR)
- Option to restore the virtual machine configuration during a Restore to Original VM.
- VM FLR enhancement: Support for FLR of Linux virtual machines that contain unformatted disks
- Option to override the automatic protection engine selection and manually select the VM Direct protection engine to use for the virtual machine image level restore

- FLR now supports LVM 2.0 for Linux distributions, and filters unformatted disks for Windows.
- The total number of concurrent TSDM virtual machine backups and restores that can be performed per ESXi host has increased from a maximum of 10 to a maximum of 18.
- Includes option to restore the BIOS UUID of a virtual machine that exists in the backup copy.
- Supports encryption on wire/in-flight for VADP and TSDM data movers. Once enabled, any further Dell DD Boost connections that happen use the encryption setting.
- Qualification of virtual machine image-based protection (VADP, TSDM) in Oracle Cloud VMware Solution
- Option to configure a nondefault port for the communication. The supported port range for Data Manager is from 7000 to 7020, excluding 7010 and 7011.

Virtual machine backup options

VM Consistent Backup

A VM Consistent Backup captures all the virtual machine disks simultaneously and backs up the data to storage targets to create a transactional-consistent backup. This option can be used for Windows and Linux virtual machines, and for guest operating systems that have applications other than the SQL Server.

Application Aware Backup

Application-aware full backup is an extension of VM full backup. For virtual machines with a SQL application installed, select this type to quiesce the application to perform the SQL database and transaction log backup. When this type is selected, Windows account credentials need to be provided for the virtual machine.

Credentials can be provided at the policy level and at the virtual machine asset level. When the credentials are provided at both the policy level and the virtual machine asset level, the virtual machine asset credentials override the policy credentials for that virtual machine. These credentials are required because Data Manager interacts with the guest virtual machines to install the Microsoft application agent and quiesce the application for performing application-consistent backups.

The agent also enables the application administrator to perform self-service restores by using the native Microsoft SQL Studio Management interface.

Exclusion

Select this type if there are virtual machine assets within the protection policy that needs to be excluded from data protection operations.

vStorage API for Data Protection (VADP) snapshot

Change Block Tracking (CBT) is supported. It allows backup applications to determine the delta of changes in the virtual machine since last backup and only read and transfer those changes when doing the next backup incrementally. Data Manager uses one or more virtual machine proxies for reading and transferring virtual machines' disk changes.

Any L0 backup of a virtual machine reads the entire contents of all disks and writes the same to storage using DD Boost (leveraging global deduplication). Any non-L0 backup of a CBT-enabled virtual machine will only read changes in the disks from the last backup. It will overlay those changes on a copy of the last backup to generate a new full backup (while moving only incremental changes).

Backup files are written to storage using fixed size segments (FSS) of 8K. Backup files on storage are always thick. For example, VMDK file-size on storage is equal to the size of provisioned disk.

Transparent Snapshots Data Mover (TSDM)

TSDM is a new protection mechanism in Data Manager v19.9 and later designed to replace the VMware vStorage API for Data Protection (VADP) protection mechanism for both crash-consistent and SQL application-aware virtual machine protection.

The advantages of using the TSDM protection mechanism for virtual machine data protection including:

- Eliminates the latency and performance impact on the production virtual machine during the protection policy life cycle
- Reduces the CPU, storage, and memory consumption required for backups. After the initial full backup, only incremental backups using the immediate previous snapshot will be performed
- An external VM Direct Engine is not required. The VM Direct Engine embedded with Data Manager is sufficient
- Automatic scaling

Starting with Data Manager 19.13, encrypted virtual machines (vmCrypt) can be protected and recovered using TSDM. This support is from vSphere 8.0p01 or higher with these limitations:

- Restoring encrypted virtual machines to a different vCenter server where encryption is not configured is not supported.
- Restore to original VM is not supported if a production virtual machine has a different encryption status than its backup copy.

See the [VMware Virtual Machine Protection using Transparent Snapshots White Paper](#) for more details.

Transport modes

Data Manager support HotAdd and NBD transport modes, the transport mode that is selected when adding the VM Direct Engine appliance (Hot-Add, Network Block Device, or the default setting Hot Add, Failback to Network Block Device). In NBD mode, the ESX/ESXi host reads data from storage and sends it across a network to the target storage.

As its name implies, this transport mode is not LAN-free, unlike SAN transport.

HotAdd is a VMware feature where devices can be added "hot" while a virtual machine is running. Besides SCSI disk, virtual machines can add additional CPUs and memory capacity. If backup software runs in a virtual appliance, it can take a snapshot and create a linked clone of the target virtual machine and then attach and read the linked clone's

virtual disks for backup. This involves a SCSI HotAdd on the ESXi host where the target virtual machine and backup proxy are running. Virtual disks of the linked clone are Hot Added to the backup proxy. The target virtual machine continues to run during backup.

vCenter authentication

At the root level of the vCenter, set up a separate vCenter user account that is strictly dedicated for use with Data Manager and the VM Direct protection engine. Use of a generic user account such as "Administrator" might make future troubleshooting efforts difficult because it might not be clear which "Administrator" actions are interfacing, or communicating, with Data Manager. Using a separate vCenter user account ensures maximum clarity if it becomes necessary to examine vCenter logs.

Before using the vCenter user account with Data Manager, or before using the Single Sign-on (SSO) admin user with the VM Direct appliance, the user must be an administrator on the vCenter root node. Users who inherit permissions from group roles are not valid.

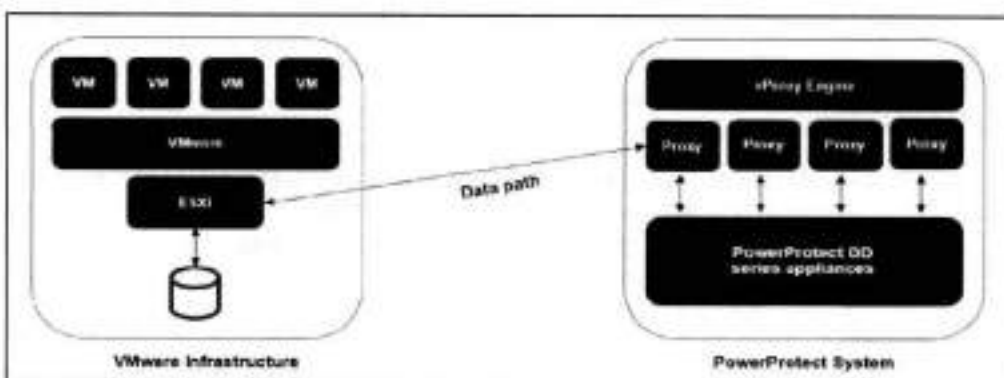
Note: ESX 6.0 U2 and above are supported for VMware Protection. By default, Data Manager enforces SSL certificates during communication with vCenter Server. If a certificate appears, click **Verify** to accept the certificate. Do not disable certificate enforcement.

VM Direct protection engine

Direct Engine is deployed in the vSphere environment to perform virtual machine snapshot backups. It improves performance and reduces network bandwidth utilization by using PowerProtect DD source-side deduplication.

Internal VM Direct Protection Engine

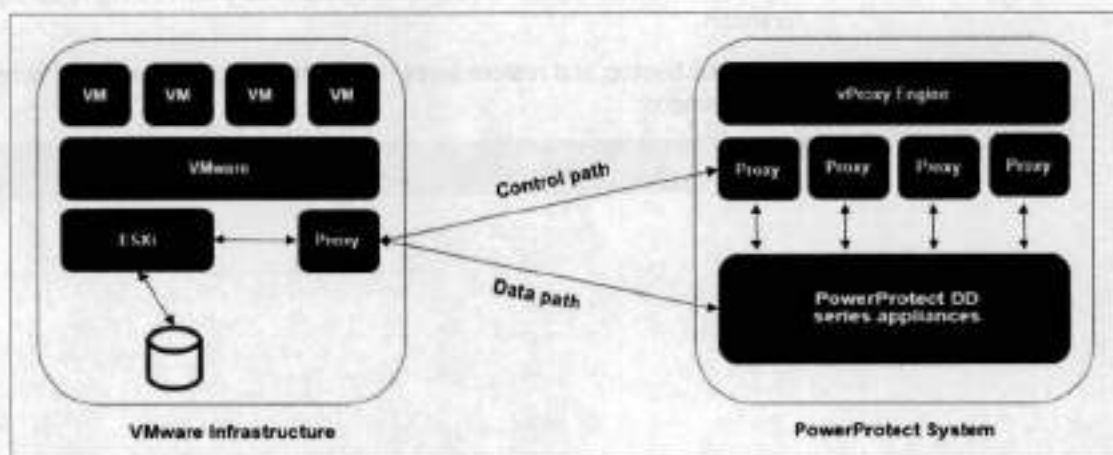
- For small scale environments
- Does not require concurrent backups
- Uses Network Block Device transport mode
- Data is transferred over network to ESXi server hosting the proxy
- Proxy gets data from its ESXi host and writes to storage



External VM Direct Protection Engine

- For larger scale environments

- Requires concurrent data protection operations
- Can also use NDB but preferred method is to use HotAdd transport mode for better performance
- Proxy attaches itself to virtual machine disk snapshot to be backed-up
- Proxy reads data from attached disk and writes to storage



Administrators can go to **Infrastructure > Protection Engines** to open the protection engines window to view statistics for the virtual machine proxy engine, manage and monitor virtual machine proxies, and add an external VM Direct Engine to facilitate data movement.



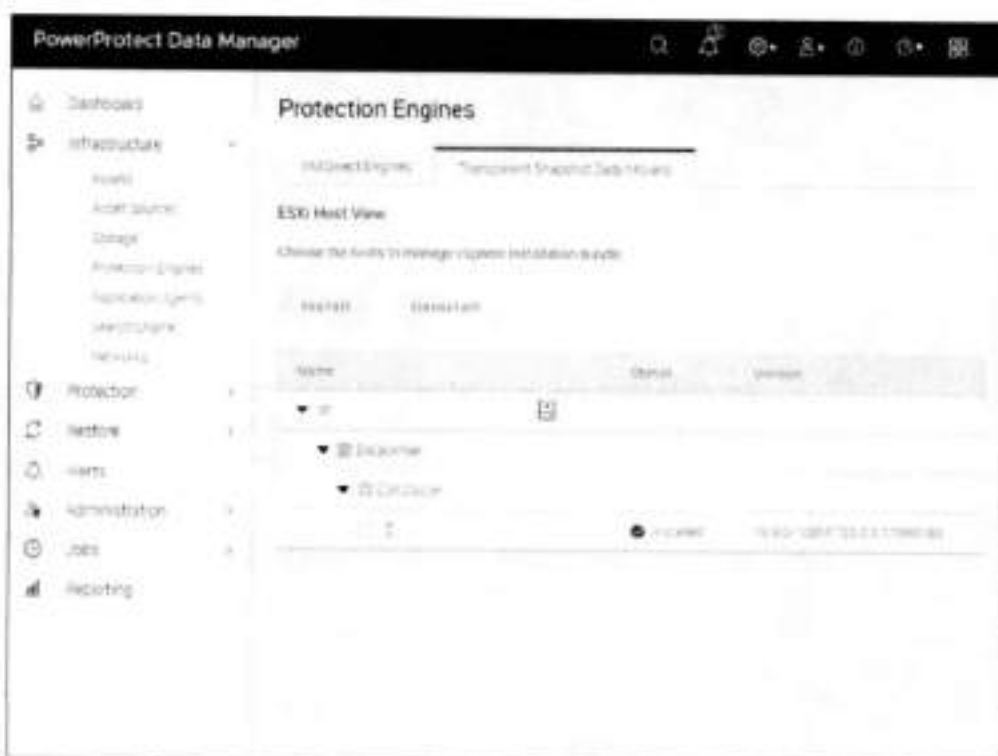
Note: When an additional VM Direct Engine is deployed and registered, Data Manager uses this appliance instead of the embedded VM Direct Engine. If an additional VM Direct Engine is not available, the embedded VM Direct Engine is used to ensure that backups are completed successfully.

VM Direct Engine (TSDM)

Data Manager comes bundled with an embedded VM Direct Engine. The embedded VM Direct Engine is sufficient for virtual machine crash-consistent and SQL application-aware protection policies that use the TSDM protection mechanism.

TSDM protection features the following improvements over VADP:

- Protect your virtual machine data using a new, optimized data path to improve overall virtual machine protection performance.
- The VM Direct Engine embedded with Data Manager can handle the virtual machine traffic. No external VM Direct Engine is required.
- Near zero latency impact to your virtual machines or environment during snapshot creation.
- Improved backup and restore times, lower infrastructure costs, and simplified management.



For more information about TSDM, see the following documents:

- [PowerProtect Data Manager—Virtual Machine User Guide on Dell Support at PowerProtect Data Manager Info Hub: Product Documents and Information](#)
- [VMware Virtual Machine Protection using Transparent Snapshots white paper](#)

Deployment requirements

Network requirements

The following table lists the key network requirements for protecting VMware virtual machines using Data Manager:

Table 1. Key network requirements

Description	Communication	Port
SSH communication	Bi-directional communication between the SSH	22 TCP
VM proxy agent management	Outbound	9613
VM proxy agent on protected guest virtual machine	Inbound	9613
vCenter communication	Bi-directional	443

- Ensure Data Manager server time is synchronized with the ESXi host system time. It is critical to PowerProtect operation that the Data Manager server time matches the systems that it interfaces with. Dell recommends that the ESXi host, and all the systems that the ESXi host interfaces with, be configured to use NTP server.
- Use Fully Qualified Domain Names (FQDNs) where possible.
- Ensure that forward and reverse DNS lookups work for each host in the protection.

VMware infrastructure requirements

To deploy Data Manager in a VMware environment, a minimum of vSphere version 6.0 is required. For more information, see *the PowerProtect Data Manager—Virtual Machine User* on Dell Support at [PowerProtect Data Manager Info Hub: Product Documents and Information Guide](#).

Supported hardware/software platform may get changed in subsequent releases. For the latest product information, see the Dell Technologies online interoperability portal: <https://elabnavigator.dell.com/eln/modernHomeDataProtection>.

TSDM requirements

The following software is required to automatically enable use of TSDM for virtual machine data protection operations.

Table 2. Software requirements for TSDM

Software required	Version supported	Notes
vCenter Server	7.0 U3	vCenter and ESXi 7.0 U3 is the minimum version that is required to use TSDM
ESXi Server	7.0 U3	
PowerProtect Data Manager	19.9 and later	

For more information, see *the PowerProtect Data Manager—Virtual Machine User Guide* on Dell Support at [PowerProtect Data Manager Info Hub: Product Documents and Information](#).

Data Manager protection policy options for virtual machine protection

A protection policy enables you to select a specific group of assets that needs to be backed up. A virtual machine protection policy can be created using Data Manager UI.

Starting with Data Manager 19.14, TSDM is the default policy. Users also have an option to migrate to VADP.

Note: Any policy created as VADP before 19.14 will remain as it is after migration to 19.14. Users will have an option to migrate the policy to TSDM later on.

Dell Technologies recommends distributing virtual machine asset protection workloads over multiple ESXi hosts so that they do not exceed the ESXi NBD session limit. If the limit is reached, workloads can be managed by deploying an external VM Direct Engine on the host/cluster using Hot Add transport mode. Data Manager provides the following backup options to back up virtual machines.

Crash Consistent



Application Aware



Policy Asset View options: View by Host and View by Asset Table

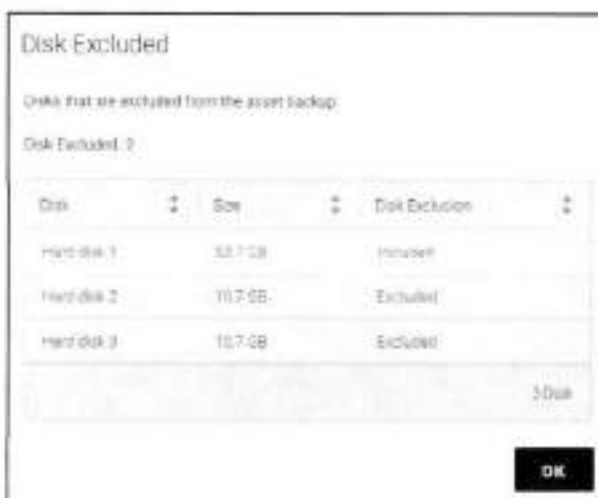


Excluding VM disks from protection

Data Manager provides the option to exclude VMDKs from virtual machine data protection during protection policy creation or by editing the existing virtual machine protection policy.

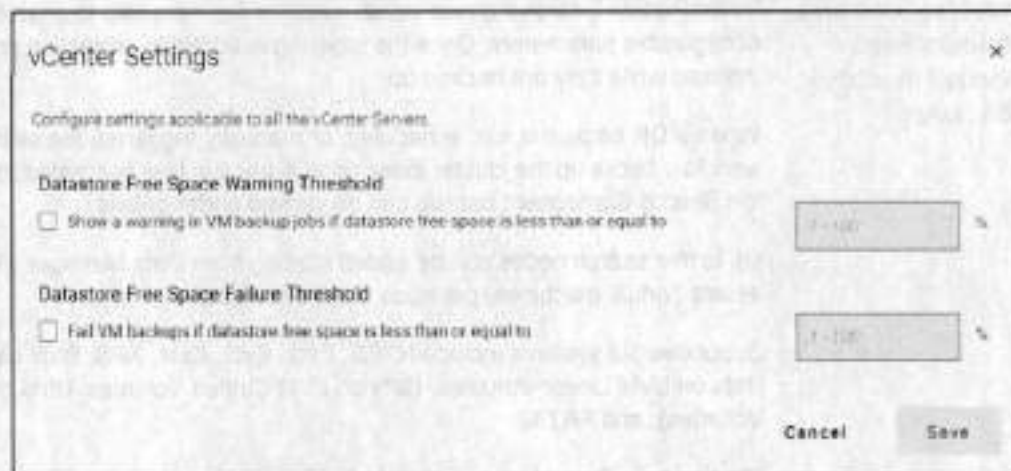


The backup copy in the following figure (in the **Recovery > Assets** section) also provides details about the disks that are excluded and included.



Datstore free space threshold settings

Datstore free space threshold settings enable warning and failure thresholds for backup and restore operations when the available space on the datstore is low.



Backup optimization modes

- Capacity to optimize for storage consumption
- Performance to optimize for speed

Option to exclude memory and swap files

The following figure shows the option to improve backup performance by excluding memory and swap files (C:\swapfile.sys, C:\pagefile.sys, and C:\hiberfil.sys) from Microsoft Windows virtual machine backups.



PowerProtect Search Engine

Search and restore files and folders from virtual machine backup

PowerProtect Search Engine is installed by default when Data Manager is installed. The PowerProtect Search indexes virtual machine file metadata to enable searches based on configurable parameters. Once the indexing is added to protection policies, the assets are indexed while they are backed up.

When a DR backup is run, scheduled, or manually triggered, the search cluster backup workflow backs up the cluster index data. A backup task is created, the individual status of the Search Component backup can be viewed under details.

Up to five search nodes can be added starting from Data Manager v19.5, and 1,000 assets (virtual machines) per node can be indexed.

Supported file systems include NTFS, Ext2, Ext3, Ext4, XFS, Btrfs (Btrfs Sub Volume, Btrfs on LVM Linear Volumes, Btrfs on LVM Striped Volumes, Btrfs on LVM Mirrored Volumes), and FAT32.

Starting from Data Manager 19.6, Indexing status at copy level for all virtual machine backups can be determined as shown here.



Success: All disks in a backup are successfully indexed.

Partial Success: Only some of the disks are indexed and others are not indexed.

In-Progress: Indexing is in progress for a backup.

Failed: Failed to index the backup.

N/A: Indexing not turned on in policy, indexing purged owing to global expiration job, unsupported File Systems.

Prerequisites:

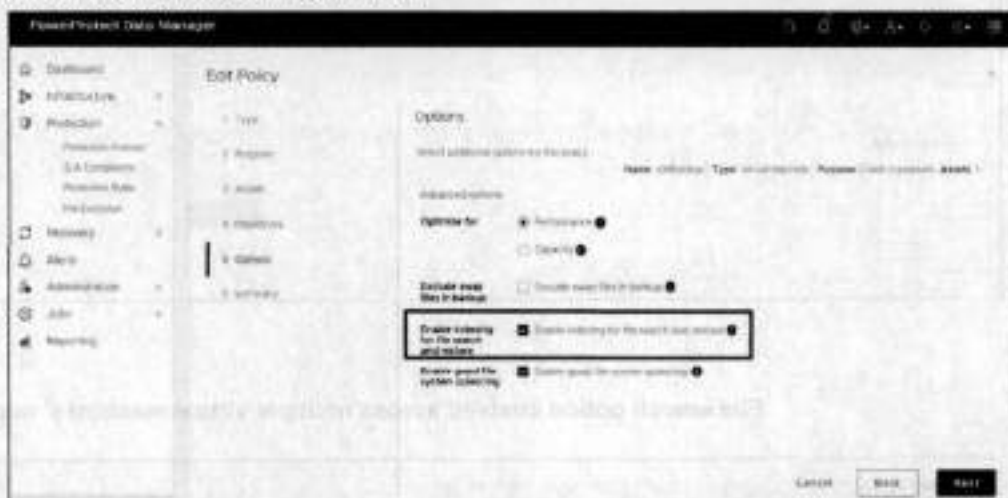
- A vCenter datastore has been configured.
- Data Manager has discovered the networks for the vCenter Server.

Search Engine deployed on Data Manager



Note: This option allows you to select another vCenter that is registered with Data Manager for Search Engine deployment. The default priority is for the vCenter where Data Manager is installed.

Option to enable indexing for file search and restore in the virtual machine protection policy configuration



File search option enabled across multiple virtual machines' copies and multiple virtual machines

Starting from Data Manager v19.8, users have the File Search option to select multiple files across multiple copies and virtual machines for both Restore to Original and Restore to Alternate options.

Supported use cases:

- Searching for files from its respective copies across different virtual machines of different kinds (Windows and Linux)
- Searching for files from its respective backup copies for similar kinds of virtual machines

- Searching for files and trying to restore the file from its various backup copies at different Point-InTime (PIT)

File search option enabled across multiple virtual machines



File search option enabled across multiple virtual machines' copies





The following figure shows options for performing a file search. Table 3 describes the file search criteria.



Table 3. File search criteria

Criteria	Description
File Name	Name of the file or folder
File Type	For example - .txt, .xlsx, .pdf
Folder Path	Path of the folder
Size	Size of the file or folder
vCenter Name	vCenter where virtual machine is hosted
VM Name	Name of the virtual machine
Backup Date	Specific backup date of the file or folder
Data Modified	File or Folder modified date
Date Created	File or Folder creation date
Guest OS	virtual machine operating system information

Search Engine node deletion from UI

Starting from Data Manager v19.8, the deletion of a Search Engine node is supported from a multinode search cluster in the Data Manager UI.

An operational node can be deleted from a search cluster to decrease cluster capacity if the space is no longer required. Nodes that could not be successfully added to the Search Engine can be redeployed or deleted.

When an operational node is deleted, Data Manager has the option to move the index data to the remaining nodes for avoiding any data loss.



Supported use cases:

Use case 1: User would like to delete the search node and move indexing data to remaining node.



2. Select the node from the decommissioning ESX host and perform delete node without data loss.



When a node is deleted, the operation is triggered and a new job is created, which you can view in the **Jobs > System Jobs** window to track its progress.

Note: Only search and indexing operations are available while the node deletion is in progress.

For the steps to delete operational nodes from a search cluster and to redeploy or delete failed nodes from a search cluster, see the *PowerProtect Data Manager Administration and User Guide* on Dell Support at [PowerProtect Data Manager Info Hub: Product Documents and Information](#).

Search Engine node deletion using CLI

Users can delete a node using CLI with '*infra node management*' which deletes the node session and removes the node from vCenter.

```
bin/infra nodegmt delete -node_id=36790321-5b96-4de7-bec3-9e1ad59d3a5e
```

Note: Recommended only when user is trying to delete all the nodes or a failed node as this would not retain any search and indexing data and cause data loss.

Search Engine node deletion using API

Users can delete a node using the API without data loss, where the data is moved to rest of the nodes and then the node is deleted.

```
POST https://0.0.0.0:8443/api/v2/search-clusters/{cluster-id}/nodes/{node-id}/management
```


```
Payload : {"nodeId": "xxxxx-xxx-xxx-xxx-xxxxxx", "operation": "DECOMMISSION"}
```


Virtual machine consistent backup workflows

Virtual machine consistent backup configuration workflow

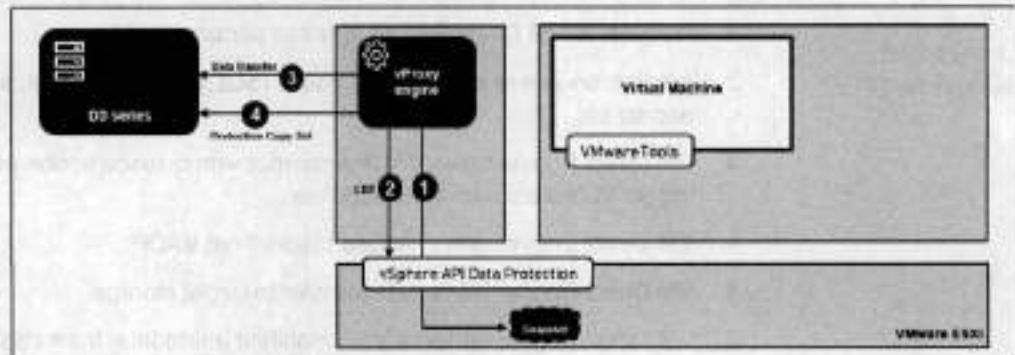
The configuration workflow of VM consistent backups is as follows:

1. User creates protection policy using PowerProtect UI.
2. Data Manager creates Boost user and storage-unit on target storage.
3. Data Manager adds protection schedule to its own scheduler.



Virtual machine consistent backup protection workflow

The protection workflow of virtual machine consistent backups is as follows:

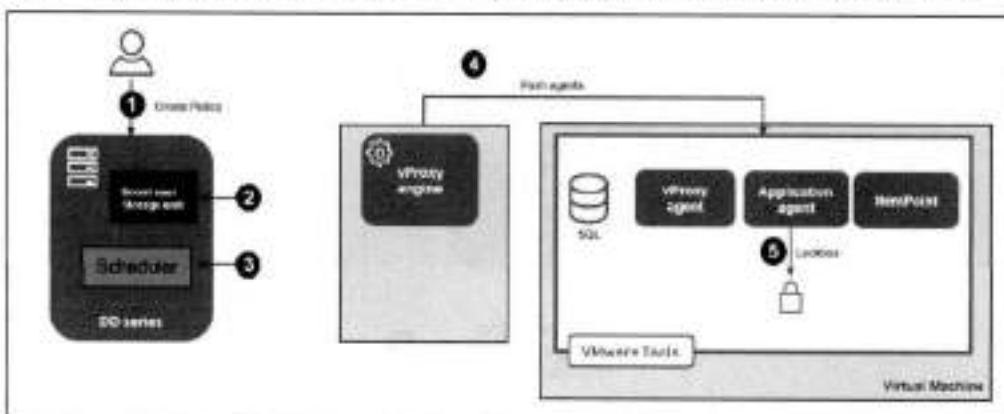


1. Virtual machine proxy takes VADP snapshot.
2. Virtual machine proxy gets changed blocks from VADP.
3. Virtual machine proxy starts data transfer to target storage.
4. Data Manager creates virtual machine protection copy set (PCS) based on the backup results.

Application-aware backup configuration workflow

The configuration workflow of application-aware backups is as follows:

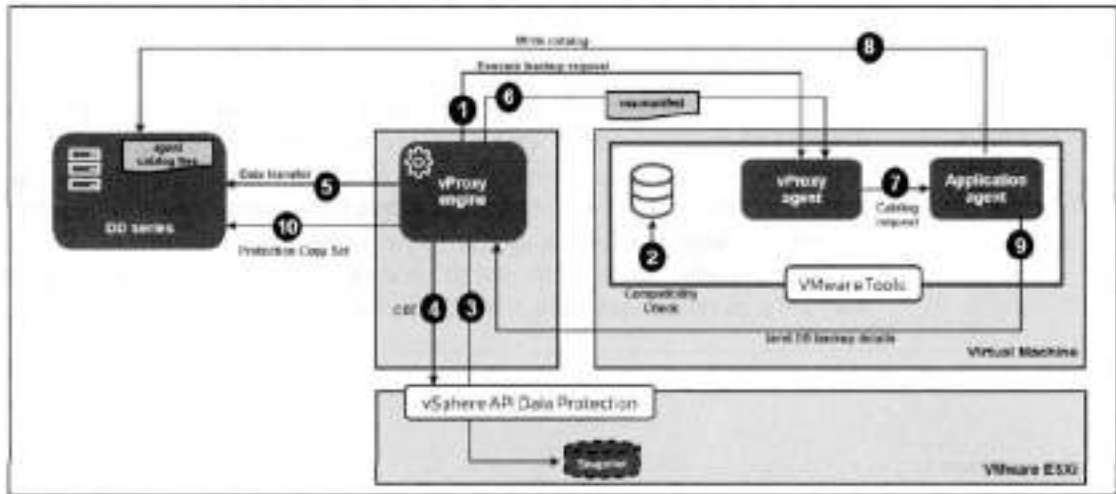
1. User creates protection policy using Data Manager.
2. Data Manager creates Boost user and storage-unit on target storage.
3. Data Manager adds protection schedule to its own scheduler.
4. VM Direct Engine pushes agent using guest operating system credentials.
 - Microsoft application agent
 - Dell VM Direct Engine agent
 - Dell ItemPoint
5. Application agent configures lockbox with credentials on SQL host (using ADM).



Application-aware backup protection workflow (FULL)

The protection workflow of application-aware database backups is as follows:

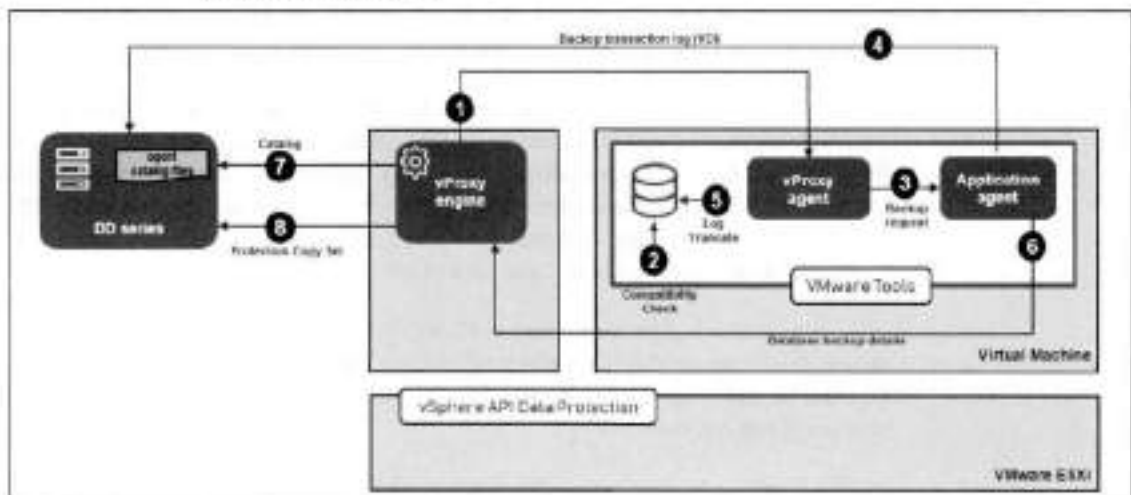
1. Request to VM Direct Engine agent to perform backup
2. Application-aware compatibility check (SQL permission, SQL status, VSS status, and so on)
3. VM Direct Engine takes VADP snapshot with quiesce option which will internally trigger VMware's own VSS workflow
4. VM Direct Engine gets changed blocks from VADP
5. VM Direct Engine starts data transfer to target storage
6. VM Direct Engine retrieves VSS manifest (metadata) from vSphere using VADP API and uploads it to the guest virtual machine
7. VM Direct Engine tells Microsoft app agent to catalog the backup
8. App agent parses VSS manifest and catalogs databases quiesced during step 3 under its own directory structure on storage
9. App agent provides database backup details, including discovered SQL assets, to VM Direct Engine
10. Data Manager creates virtual machine protection copy set (PCSt) and corresponding SQL PCS based on the backup results



Application-aware backup workflow (log)

The protection workflow of application-aware log backups is as follows:

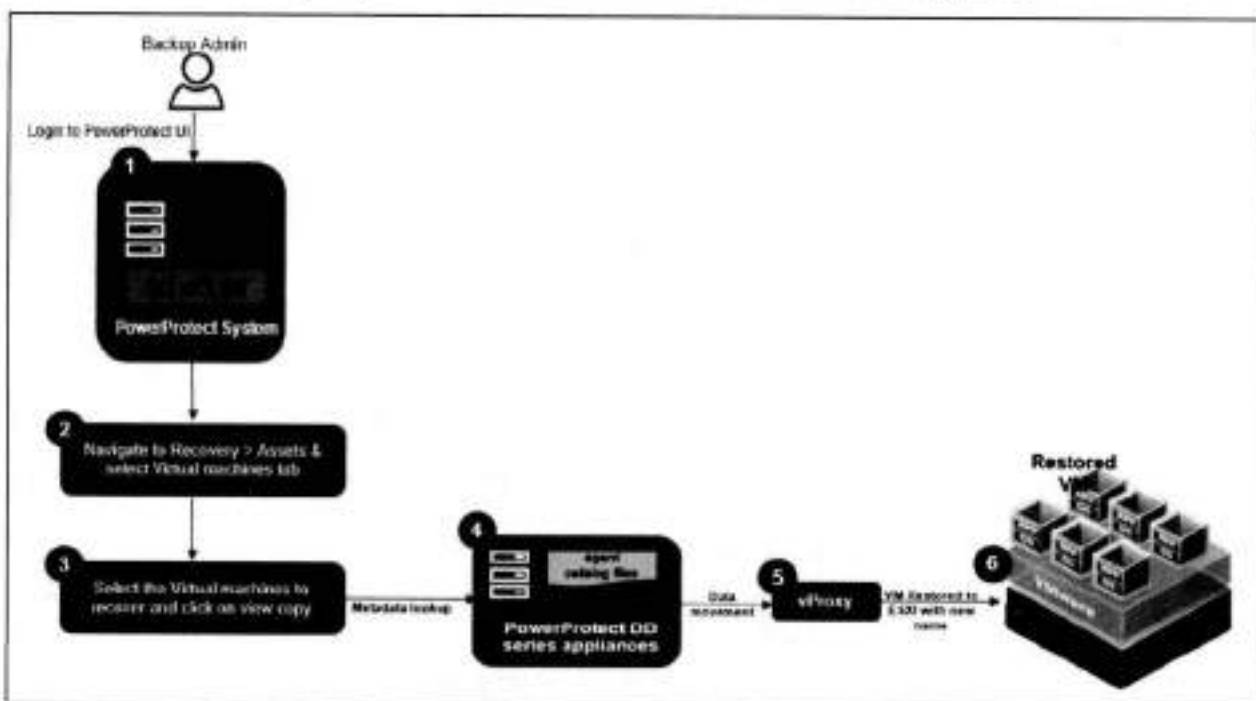
1. Request to VM Direct Engine to perform backup.
2. Application-aware compatibility check (SQL permission, SQL status, VSS status, and so on).
3. VM Direct Engine asks Microsoft App agent to perform transaction log backup.
4. Microsoft App agent will serially back up each database transaction log (using VDI) directly to target storage.
5. SQL Server truncates logs.
6. App agent provides database backup details, including discovered SQL assets, to VM Direct Engine.
7. VM Direct Engine parses VSS manifest and catalogs files and transaction logs.
8. Data Manager creates virtual machine PCS and its corresponding SQL PCS based on the backup results.



Data Manager use cases for virtual machine recovery

Use cases

Restore to new: Create a new virtual machine using a copy of the original virtual machine backup and can be restored on the vCenter server using a different name. The virtual machine can also be restored on an alternate vCenter server if Data Manager has previously discovered it. Instant access allows the virtual machine to be created and powered on while temporarily accessing the vmdk from PowerProtect. The virtual machine becomes available for use as soon as it is powered on. The following flow diagram provides an overview of the Restore to new recovery operation.



Restore to original: Recover virtual machine backup to its original location on the same vCenter. Roll back the virtual machine that was protected to an earlier point in time. Unlike Restore to New, there are no options to be selected. A dialog box will appear, requesting confirmation to restore this virtual machine.

Live virtual machine: Creates a new virtual machine directly from the original virtual machine backup for the purposes of instant backup validation and recovery of individual files. This process does not copy or move any data from storage to the production datastore. VMware administrator can vMotion the virtual machine manually. The live virtual machine is initially available for 7 days. Monitor and manage the live virtual machine recovery from the Instant Access menu.

File level restore: File level restore allows recovering individual files from backups of virtual machines or VMDKs performed in Data Manager to a primary or secondary vCenter server. File-level restore is only supported for the following platforms and operating system versions.

VMDK Restore to Alternate VM: Starting with Data Manager 19.13, individual VMDKs can be restored from a backup copy to any existing virtual machine including the original

virtual machine. VMDKs restore to alternate adds to the existing virtual machine. This feature is also known as Granular VMDK Restore to Alternate or Recover VMDKs.

DD Boost compressed restore

DD Boost compressed restore improves backup read performance by using data compression techniques. This option enables DD Boost compressed restore to improve DD Boost backup read performance. Compression reduces restore times but increases CPU usage on both systems.

Restore VM sessions data is compressed on the PowerProtect DD series appliance before transmission to VM Direct Engine where it is uncompressed. This functionality results in reduced network load that provides up to 50 percent reduction (based on in-house testing) in restore times for bandwidth-limited setups.

Data Manager provides an option to enable the DD Boost compressed restore from the Data Manager UI and from restore VM sessions API for VM Direct Engine.

The following screen shows the option to Enable DD Boost Compression from Data Manager UI during restore:

The screenshot shows a 'Restore' configuration screen with three main sections: 'Purpose', 'Restore Type', and 'Restore Options'. The 'Restore Options' section is highlighted with a red box and contains the following items:

- Enable DD Boost Compression
- Restore VM Tags
- Restore Storage Policies
- Enable DD Boost Compression

The following screen shows the option to Enable DD Boost Compression from restore VM sessions API for VM Direct Engine:

The screenshot shows the 'Restore VM Sessions' configuration screen. The 'Enable DD Boost Compression' checkbox is highlighted with a red box. The checkbox is currently unchecked. The text next to it reads: 'DD Boost compressed restore improves backup read performance by using data compression techniques. This option enables DD Boost compressed restore to improve DD Boost backup read performance. Compression reduces restore times but increases CPU usage on both systems.'

When set to false, the application agent does not enable DD Boost compressed restore. When set to true, the application agent enables DD Boost compressed restore.

This option can be controlled by using the API `enableCompressedRestore` property in a POST `/api/v2/restored-copies` API operation.

For more information, see the recovery and reuse management section of the [PowerProtect Data Manager REST API documentation](#).

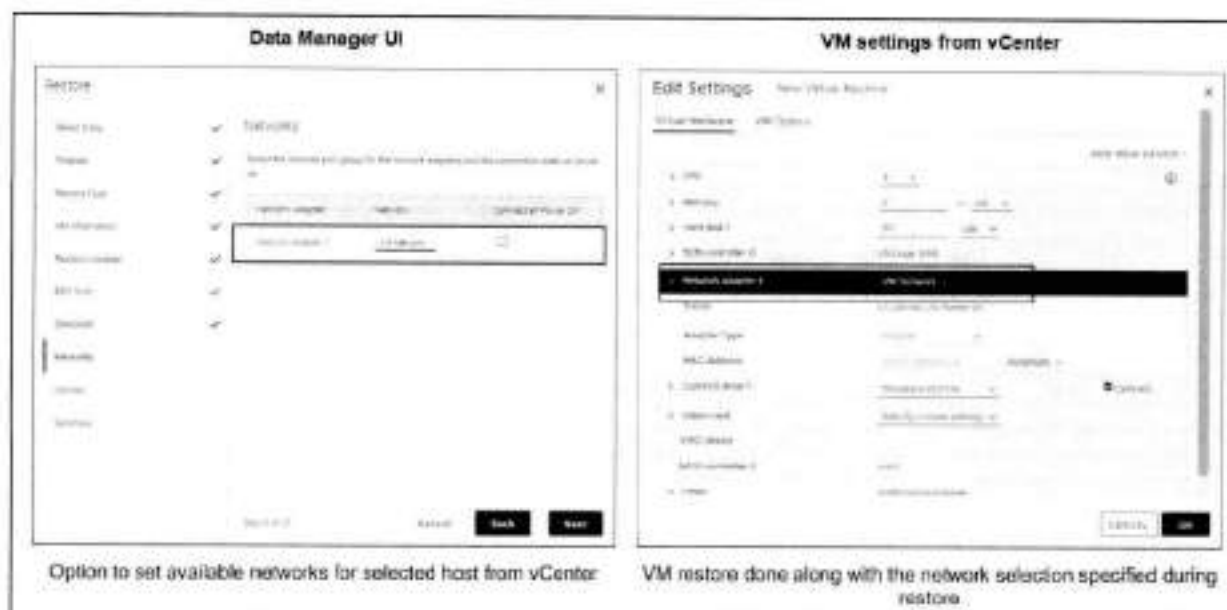
DD Boost compressed restore applies to the following use cases:

- Direct restore to ESXi
- Create and restore to new virtual machine
- Restore to original folder and overwrite original files
- Restore individual disk

Note: DD Boost compressed restore does not apply to NFS-based restore (Instant restore).

Change virtual machine network settings during restore

The network settings of a virtual machine can be changed during the virtual machine restore. This allows you to select a different network for an adapter if the network used by the original virtual machine is no longer available. The initial "Power On" connection status of any network adapter can also be modified. This only applies to virtual machines that have been backed up with Data Manager v19.9.



Note: When the selected network is not available during restore, the restore would fail.

Restore virtual machine configuration during a Restore

Starting with Data Manager version 19.10, you can restore the virtual machine configuration during a Restore to Original VM. Select Restore VM Configuration to use the .vmx file to restore the virtual machine configuration that existed at the time of the backup. If there were changes to the virtual machine disk configuration, you cannot clear this option.



Override User Account Control (UAC)

An Override User Account Control (UAC) option is added for Windows and Linux file level restores. On Windows, the local user must be part of the Administrators group. On Linux, this account requires sudo access and must be able to run vrfcopy without being prompted for a password.



Override automatic VM Direct protection engine selection

Starting with Data Manager version 19.11, the UI provides an option to override the automatic protection engine selection and manually select the VM Direct protection engine to use for the virtual machine image level restore.



Manage and monitor instant access sessions

In the Instant Access Sessions window, you can change the status of a virtual machine restore to new or instant access virtual machine restore (for example, by extending the availability period or deleting an instant access virtual machine) and monitor vMotion events.

Note: The instant access sessions used by a SQL application-aware self-service restore are displayed in the Data Manager UI, but management is disabled. Use the SQL application-aware self-service restore UI to manage these sessions.

When the Jobs window indicates that a recovery has completed successfully, go to **Restore > Running Sessions** to access the **Running Sessions** window. In this window, you can monitor and manage all exported copies created from storage.

An active restore session with a state of **Mounting** indicates that the restore is still in progress. Once the state changes to **Mounted**, the restore is complete and the instant access virtual machine is ready. When the checkbox next to the session is selected, one option can be chosen from three options, as shown in the following figure.



Orchestrate the restore of assets through Restore Plan

Starting with Data Manager 19.14, the UI includes a Restore Plan option that can be used to orchestrate the restore of massive assets. Users can create a restore plan for massive assets and then run it whenever needed or based on a schedule.

A restore plan can be reused multiple times. Users can view, create, edit, and delete a restore plan.

A restore plan can be applied to **All Assets** or **Select Resource Group**. Only scope users can run a restore plan for assets that the user can access



Restore groups

A restore plan consists of one or multiple restore groups. Restore groups contain everything that is required for restoring assets:

- Asset selector—Select assets to restore.
- Copy selector—Select copy for each asset to restore.
- Target location—Set a location for the target asset.
- Configurations—Support all the configurations in the legacy copy restore.
- Priority—Set priority in which restore groups are run.

Data Manager use cases for virtual machine recovery

Add Restore Plan

Information | Description

Name for Restore Plan and select a Scope

Name:

Description:

Scope: All Assets Select Resource Group

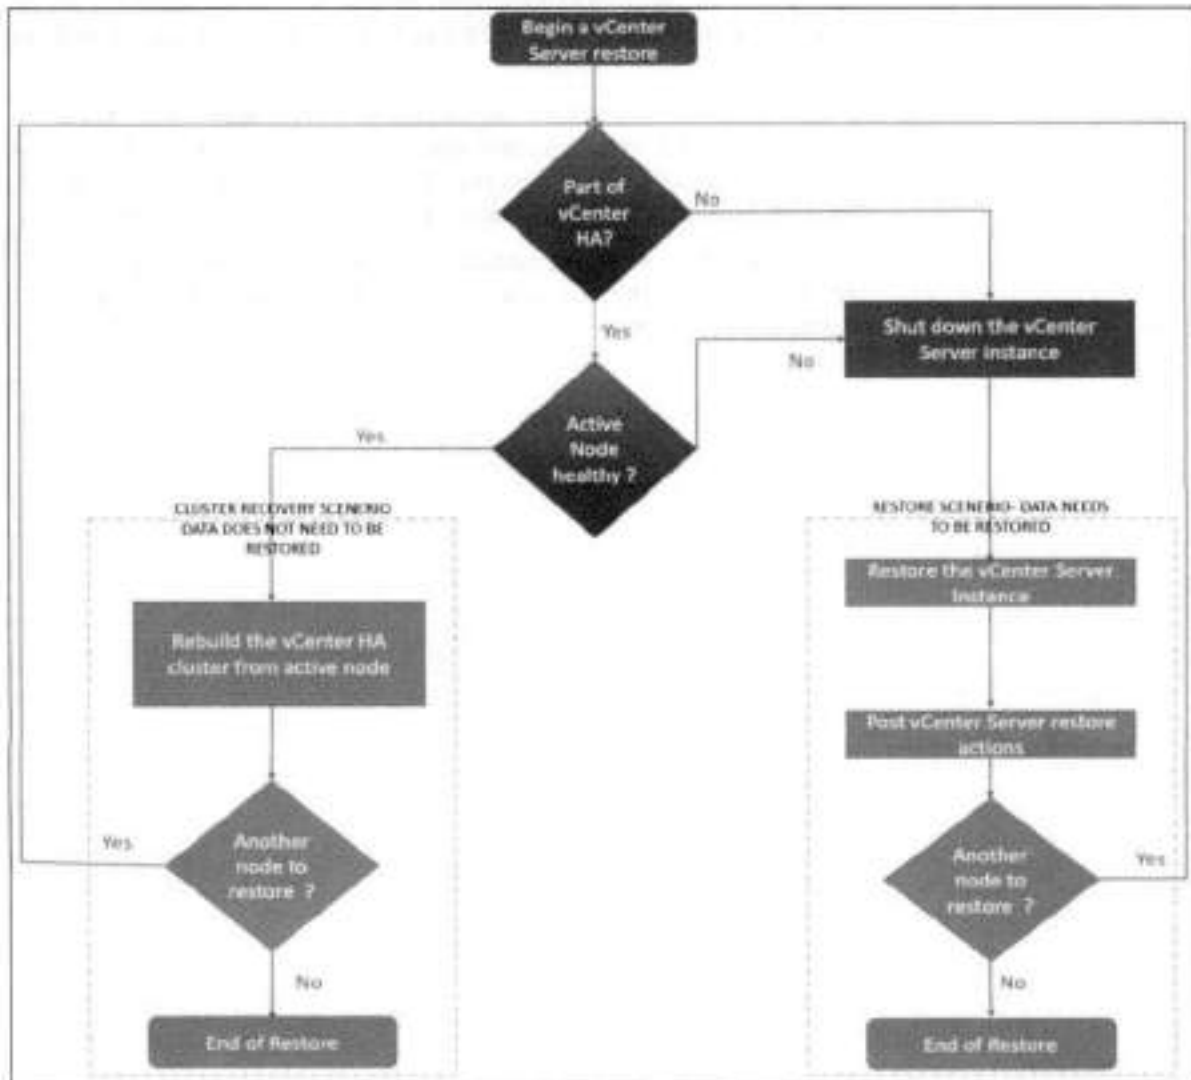
<input type="checkbox"/>	Details	Resource Group Name	Description	Assets
<input type="checkbox"/>	TS	ResourceGroup1		1
<input type="checkbox"/>	GS	ResourceGroup2		1

Note:

- Restore Plan only supports VM in 19.14.
 - Max number of assets in one restore group is 500.
 - Now allowed to set new names for the restored assets.
-

Disaster recovery

Data Manager supports protecting vCenter 6.5 deployments using VM Direct Engine appliance. It will be useful in case disaster or vCenter down scenario to recover the virtualized environments. Following are the recommendations and best practices when planning a vCenter virtual machine or its component virtual machines backup.



- Schedule the backup of the vCenter server when the load on the vCenter server is low, such as during off-hours, to minimize the impact of vCenter virtual machine snapshot creation and snapshot commit processing overhead
- Ensure that there are no underlying storage problems that might result in long stun times
- Keep the vCenter virtual machine and all its component virtual machines in one single isolated NetWorker group/policy. This is to ensure that the backup times of all vCenter server component virtual machines are as close to each other as possible

- If using one or more external Platform Services Controllers, it is recommended to have one dedicated VM Direct Engine associated to the workflow for the entire vCenter server virtual machines backup. This will ensure that the backup times of all vCenter Server component virtual machines are as close to each other as possible
- Set the maximum HotAdd session limit of the dedicated VM Direct Engine to an appropriate number to avoid queuing of backups. It is recommended to set the maximum HotAdd session limit to 25 and the maximum NBD session limit to 0 (zero)
- Ensure that the backup start time of the vCenter Server does not overlap with any operations for other protected virtual machines being managed by this vCenter. Doing so ensures that there is no impact on other protected virtual machines during snapshot creation and snapshot commit of the vCenter virtual machine
- If the vCenter Server and Platform Services Controller instances fail simultaneously, the Platform Services Controller must be restored first and then the vCenter Server instances

Technical support and resources

Dell.com/support is focused on meeting customer needs with proven services and support.

The following PowerProtect Data Manager documentation is on Dell Support at [PowerProtect Data Manager Info Hub: Product Documents and Information](#):

- [PowerProtect Data Manager - Virtual Machine User Guide](#)
- [PowerProtect Data Manager Administration and User Guide](#)
- [PowerProtect Data Manager Deployment Guide](#)

The [Data Protection Info Hub](#) provides expertise that helps to ensure customer success with Dell data protection products.

Dell EMC PowerProtect Data Manager Protecting VMware Tanzu Kubernetes Clusters

December 2021

H18682.1

White Paper

Abstract

VMware vSphere with Tanzu transforms vSphere clusters into a platform on which Kubernetes workloads can be run directly on VMware ESXi hosts and can create Kubernetes clusters within dedicated namespaces. This document describes the architecture and configuration of VMware Tanzu Kubernetes Grid clusters with Dell EMC PowerProtect Data Manager and explains how the VMware vSphere with Tanzu workloads are protected.

Dell Technologies

Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA December 2021 H18682.1.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Contents

Executive summary	4
Introduction and components	5
Prerequisites for enabling vSphere with Tanzu	7
Architecture	8
Configuring VMware vSphere with Tanzu	12
Configuring PowerProtect Data Manager to protect Tanzu Kubernetes workloads	21
Technical support and resources	33

Executive summary

Overview

Modern IT infrastructure is being transformed by Containers. Containers are similar to virtual machines but have relaxed isolation properties to share the operating system. The Container has its own filesystem, CPU, memory and process space. Agile application creation, continuous development, environmental consistency across development, application-centric management, efficient resource allocation and resource isolation are the key benefits of containers. Kubernetes is an open-source container management platform that unifies a cluster of machines into a single pool of compute resources.

VMware vSphere is the compute virtualization platform. VMware vSphere 7 rearchitected with native Kubernetes for application modernization that enable IT admins to use vCenter server to operate Kubernetes clusters through namespaces. VMware vSphere with Tanzu provides a platform for both traditional applications as well as modern applications so that both IT admins and developers can access developer-ready infrastructure, scale with simple operations.

With currently distributed container deployment, it is important to protect the workloads. Dell EMC PowerProtect Data Manager protects the workloads and ensures high availability, consistent, and reliable backup and restore for Kubernetes workload or DR situation. PowerProtect Data Manager offers centralized management, automation, multi-cloud options and advanced integration for ease and simplicity for managing workloads. PowerProtect Data Manager protects VMware Tanzu Kubernetes Grid (TKG) clusters, pods, persistent volume claims, namespaces and other resources.

Audience

This white paper is intended for customers, partners and others who want to understand how PowerProtect Data Manager software helps to protect VMware Tanzu Kubernetes Grid clusters and the workloads.

Scope

1. VMware Sphere with Tanzu - A Tanzu edition license of vSphere 7.0.1 and above
2. Enabling Workload Management using vSphere Distributed switch using HA proxy
3. PowerProtect Data Manager 19.7 and above

Revisions

Date	Description
February 2021	Initial release
December 2021	Template update

Note: This document may contain language from third-party content that is not under Dell Technologies' control and is not consistent with 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.

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#) (subject line: Feedback for document: H18682.1).

Author: Abhishek Shukla

Note: For links to other documentation for this topic, see the [Data Protection Info Hub](#).

Introduction and components

Introduction

The Cloud Native definition is an architectural philosophy for designing the applications and infrastructure. Containers provide a way to package and run the application. To run such applications, container orchestration is required. Kubernetes is an open-source container orchestrator for managing containerized workloads and services, facilitating both declarative configuration and automation. It is portable, extensible, and scalable and has a large, rapidly growing ecosystem. Kubernetes services, support, and tools are widely available, and the applications are constructed of multiple microservices that run a large number of Kubernetes pods and VMs. VMware vSphere with Tanzu helps in creating Kubernetes control plane directly on VMware ESXi by creating Kubernetes layer within ESXi that are part of the Kubernetes cluster.

Dell EMC PowerProtect Data Manager protects existing as well as new discovered workloads. It allows IT operations and backup admins to manage VMware Tanzu clusters and their protection through a single management UI and define protection policies for Kubernetes workloads from Kubernetes APIs. The policy-driven protection is defined by the Protection Policy mechanism. PowerProtect Data Manager discovers the namespaces, labels, and pods in the environment, which can be protected by providing cluster credentials. Logging, Monitoring, governance and recovery are done through PowerProtect Data Manager.

VMware vSphere with Tanzu components

VMware vSphere with Tanzu provides a platform for running Kubernetes workloads natively on VMware ESXi.

Workload

In vSphere with Tanzu, the workload is a deployed application that consists of containers running inside vSphere Pods, VMs, or both. It is an application that runs inside a Tanzu Kubernetes cluster that is deployed by Tanzu Kubernetes Grid service.

Supervisor cluster

The supervisor cluster provides the management plane on which Tanzu Kubernetes clusters are built. The Tanzu Kubernetes Grid (TKG) service is a controller manager that includes a set of controllers which is a subset of the supervisor cluster. TKG service helps in provisioning a Tanzu Kubernetes cluster.

Supervisor namespace

When Tanzu Kubernetes clusters are provisioned, a resource pool and VM folder are created in a supervisor namespace. The resource quotas and storage policy are applied to a namespace and inherited by the deployed Tanzu Kubernetes cluster. The Tanzu Kubernetes cluster control plane and worker node VMs are placed within the resource pool and VM folder.

Tanzu Kubernetes cluster

The Tanzu Kubernetes cluster is a distribution of the open-source Kubernetes container platform that is built, signed, and supported by VMware. Tanzu Kubernetes clusters are built on top of the supervisor cluster. The cluster is defined in the supervisor namespace using custom resource. It uses the open-source Photon OS from VMware and is integrated with underlying vSphere infrastructure including storage, network, and authentication.

vSphere pod

A vSphere pod is a VM with a small footprint that runs one or more containers. It is similar to a Kubernetes pod. Each pod is sized for the workload that has explicit resource reservations for that workload. It is allocated a specific amount of storage, memory, and CPUs that are required for the workload to run.

Persistent Volume (PV) and Persistent Volume Claim (PVC)

Persistent Volume is a storage defined for the cluster that is provisioned by an administrator or dynamically provisioned using Storage Classes (SCs). It is a resource in the cluster similar to a node. PVs are volume plug-ins like volumes but have a lifecycle independent of any individual pod that uses the PV. It captures the details of the implementation of the storage that is NFS, iSCSI, or a cloud-provider-specific storage system.

A Persistent Volume Claim (PVC) is a request for storage by a user. It is like a pod. Pods consume node resources. Similarly, PVCs consume PV resources. Pods can request specific levels of resources (CPUs and memory).

Storage Class

A Storage Class is described as the type of storage that is provisioned and allowed ranges for size and IOPS. When a user creates a PVC, the storage class is specified with size in GB and number of IOPS. A storage class is used to abstract the underlying storage platform.

Custom Resource (CR)

A resource in a Kubernetes environment is an endpoint for API that stores a collection of API objects of a certain kind. A Custom Resource (CR) is an extension of the Kubernetes API that is not necessarily available in a default Kubernetes installation. It represents a customization of a particular Kubernetes installation.

PowerProtect Data Manager components

Cloud Native Data Manager

The Cloud Native Data Manager (CNDM) is in-built microservice component of PowerProtect Data Manager that communicates with the kube-apiserver of the cluster. This component is responsible for APIs for the backup and restore process.

PowerProtect controller

The PowerProtect controller is the component that is installed on the Kubernetes cluster when the cluster is discovered by PowerProtect Data Manager. The backup and restore controllers manage BackupJob CR and RestoreJob CR definitions and are responsible for the backup and restore of Persistent Volumes.

VMware Valero

VMware Valero is an open-source tool that is integrated with PowerProtect Data Manager. It is built-in and does not have to be installed separately. Valero is pushed into the Kubernetes cluster by the PowerProtect controller pod after it is in an up-and-running state via the Valero deployment object. It is responsible for the backup and restore of metadata.

vProxy (VM proxy)

The vProxy protection engine is the virtual machine data protection component within PowerProtect Data Manager. During backups, the vProxy agent creates a snapshot of virtual-machine data directly from the datastore. The snapshot is moved directly to the target storage where the backups are stored. This process uses VMware vSphere Storage API for Data Protection (VADP) which enables centralized, off-host, LAN-free backup of virtual machines.

Note: The VADP is a subset of the vSphere API that enables backup and restore applications. The snapshot-based VADP framework allows efficient, off-host, centralized backup of virtual-machine storage. After taking a snapshot to quiesce virtual disks, software can offload the backup load to the target storage.

Prerequisites for enabling vSphere with Tanzu

Introduction	To configure or run Kubernetes workloads natively on vSphere, workload management must be enabled to create a supervisor cluster where the vSphere pods run and provision the Tanzu Kubernetes cluster. There are few prerequisites for compute, network, and storage.
vSphere cluster	<p>vSphere cluster is a collection of ESXi hosts managed by vCenter server.</p> <ul style="list-style-type: none"> • To enable workload management, at least 3 ESXi hosts are required; if you are using VSAN, at least 4 ESXi hosts are required. • vSphere cluster must be configured with high-availability (HA) enabled. • vSphere cluster must be configured with Distributed Resource Scheduler (DRS) enabled and must be set to fully automated mode. • The cluster must use shared storage for vSphere HA, DRS, and for storage persistent volumes.

Networking stack

To enable workload management, the networking must be configured for the supervisor cluster. There are two ways to configure networking such as NSX-T Data Center or vSphere Distributed Switch (vDS) with an external load balancer (HA proxy):

- To use vSphere Distributed Switch (vDS) with HA proxy load balancing for the supervisor cluster, see the system requirements at <https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-C86B9028-2701-40FE-BA05-519486E010F4.html>.
- To use NSX-T Data Center networking for the supervisor cluster, see the system requirements and topologies at <https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-B1388E77-2EEC-41E2-8681-5AE549D50C77.html#GUID-B1388E77-2EEC-41E2-8681-5AE549D50C77>.

Storage policy

Storage policies are created for the datastore placement for Kubernetes control plane VMs, containers, and images. Storage policies are associated with different storage classes. Before workload management is enabled, a storage policy is created for the placement of Kubernetes control plane VMs.

- The datastore must be shared among all ESXi hosts in the cluster.
- VM storage policies must be configured and updated.

For information about storage policy creation with vSphere, see <https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-544286A2-A403-4CA5-9C73-8EFF261545E7.html#GUID-544286A2-A403-4CA5-9C73-8EFF261545E7>.

Content library

A content library consists of distributions of Tanzu Kubernetes releases in the shape of OVA templates. You can create a Local Content Library, where images are uploaded manually, or a Subscribed Content Library to pull the latest released images automatically.

Architecture

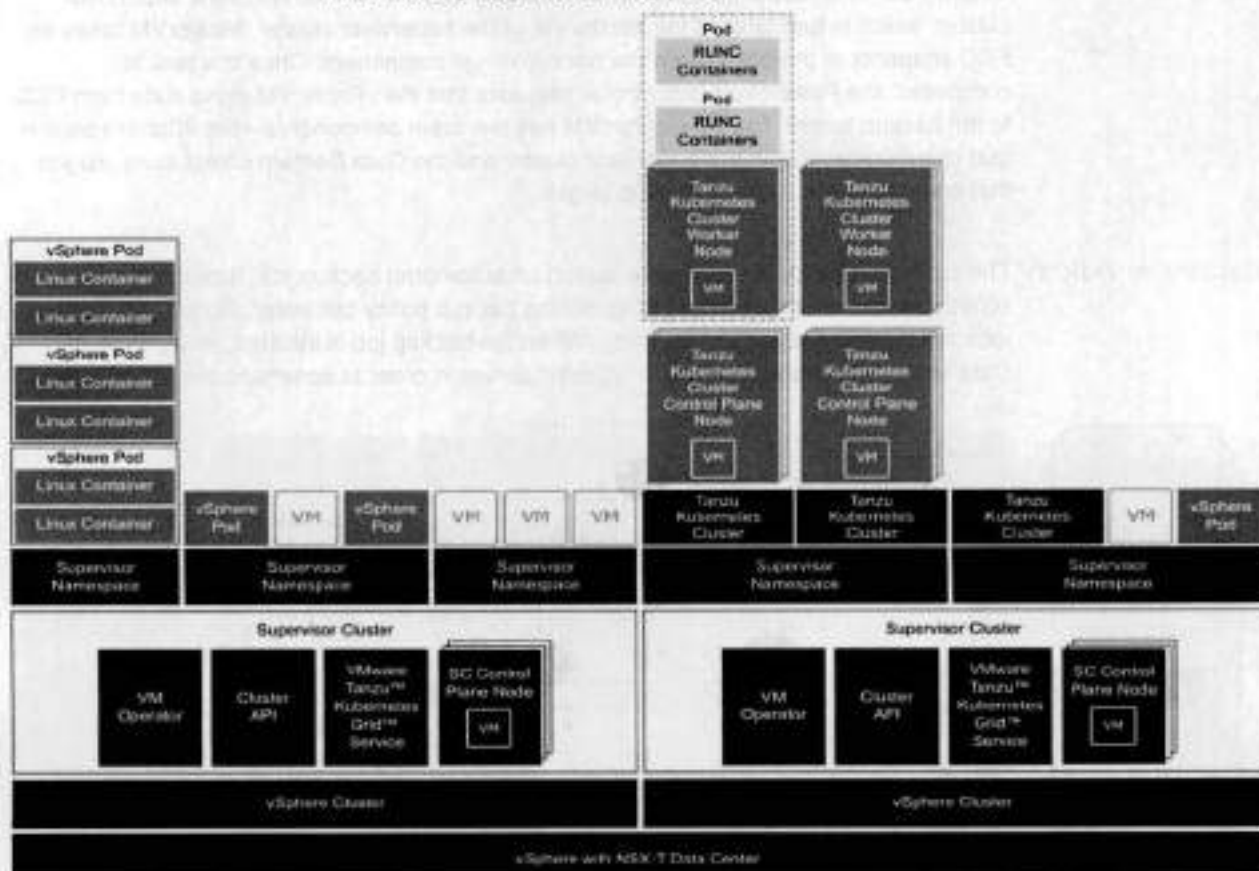
Architecture overview

PowerProtect Data Manager 19.7 introduces the ability to protect Tanzu Kubernetes cluster workloads. VMware vSphere 7U1 re-architects vSphere with native Kubernetes as its control plane. A TKG cluster is a Kubernetes cluster that runs inside the virtual machines on the supervisor layer, which allows Kubernetes to run with consistency. It is enabled via the TKG service for VMware vSphere and is upstream-compliant with open-source Kubernetes (guest cluster). The guest cluster is a Kubernetes cluster running on VMs and consists of its control plane VM, management plane VM, worker nodes, pods, and containers.

PowerProtect Data Manager protects Kubernetes workloads and ensures that the data is consistent and highly available. PowerProtect Data Manager is a virtual appliance that is deployed on an ESXi host using OVA. It is integrated with PowerProtect Data Domain Virtual Edition (DDVE) as protection target where backups are stored.

Once the cluster is discovered, it is added as a PowerProtect Data Manager asset source, and associated namespaces as assets are available to be protected. During the discovery process, PowerProtect Data Manager creates the following two namespaces in the cluster. The data is compressed and deduplicated at the source and sent to the target storage.

1. **Velero-ppdm:** Contains a Velero pod to back up metadata and stage it to the target storage in a bare metal environment. It performs PVC and metadata backup in VMware Cloud Native Storage (CNS).
2. **PowerProtect:** Contains a PowerProtect controller pod to drive Persistent Volume Claim snapshot and backup and push the backups to target storage using intermittently spawned cProxy pods.



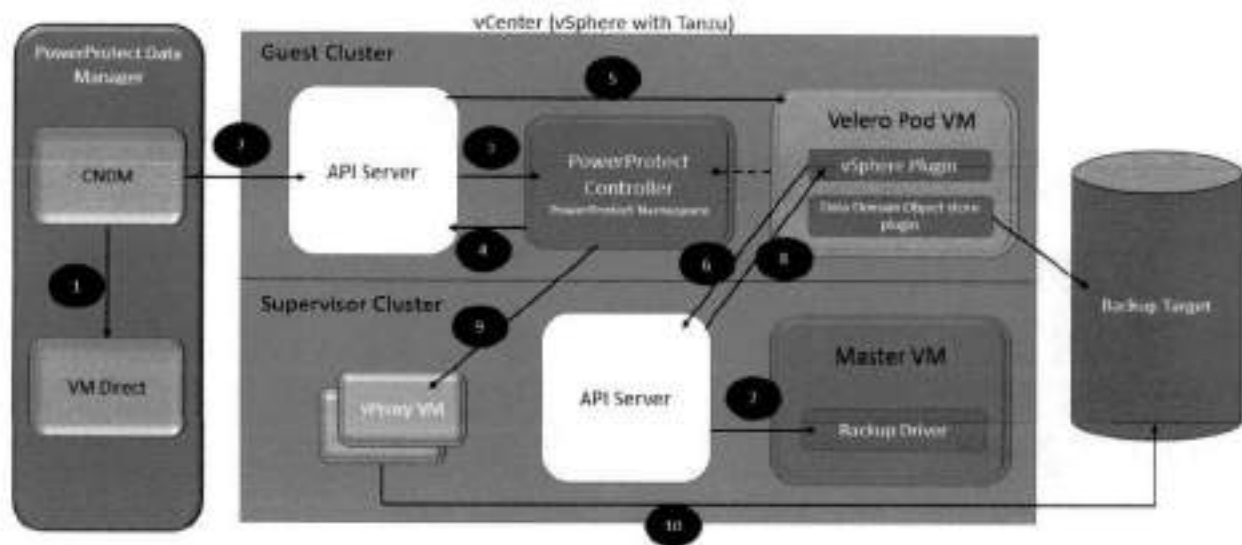
Note: The pods running in the guest clusters do not have direct access to the supervisor cluster. The persistent volumes provisioned by vSphere CSI on the guest cluster create an FCD disk in the supervisor space and are mapped to the guest cluster via paravirtual CSI driver. PowerProtect Data Manager uses internal APIs to protect paravirtual volumes.

According to Tanzu Kubernetes cluster architecture, vSphere cluster (ESXi as worker node) has supervisor clusters and guest clusters (TKG clusters). The guest clusters have their own control plane VMs, management plane, worker nodes, networking, pods, and namespaces and are isolated from each other. Supervisor clusters and guest clusters communicate via API servers. The cProxy of PowerProtect Data Manager does not have

access to the pods running on the guest clusters because it is external to the clusters; therefore, PowerProtect Data Manager does not use cProxy for the backup and restore process. However, PowerProtect Data Manager uses a vProxy based protection solution. The vProxy agent creates a snapshot of VM data directly from the datastore. The snapshot is moved directly to the target storage where the backups are stored. When the backup job is triggered, CNDM communicates with the VM direct to find and reserve a vProxy. The vProxy is created at the vCenter specifically for TKG clusters. Once the vProxy is reserved, CNDM initiates the communication with the API server of the guest cluster using Velero operator. The API server then communicates with the PowerProtect controller (PowerProtect namespace) where the backup job and Velero backup custom resources are created. It communicates with Velero PodVM.

Velero PodVM is responsible for communicating with the API server of the supervisor cluster, which in turn talks to the MasterVM of the supervisor cluster. MasterVM takes an FCD snapshot of the pods using the backup driver component. Once this task is completed, the PowerProtect controller requests that the vProxy VM move data from FCD to the backup target. The Velero PodVM has two main components—the vSphere plug-in that communicates with the supervisor cluster and the Data Domain object store plug-in that communicates with the backup target.

Backup workflow The protection policy is required to launch an automated backup job. It centrally schedules and manages timing of launching backup policy per asset. Scheduled backup jobs are triggered according to policy. When the backup job is initiated, PowerProtect Data Manager communicates with vCenter server in order to communicate with clusters.



Steps 1, 2: CNDM and VM direct are two components within PowerProtect Data Manager. When the backup is triggered, CNDM communicates with VM direct to find and reserve a vProxy at the vCenter supervisor cluster. If an existing vProxy is not found, backup jobs fail. The vProxies are created externally specifically for TKG clusters. Once the vProxy is reserved, CNDM initiates communication with the API server of the guest cluster by sending BackupJob CR using Velero operator and passing vProxy details.

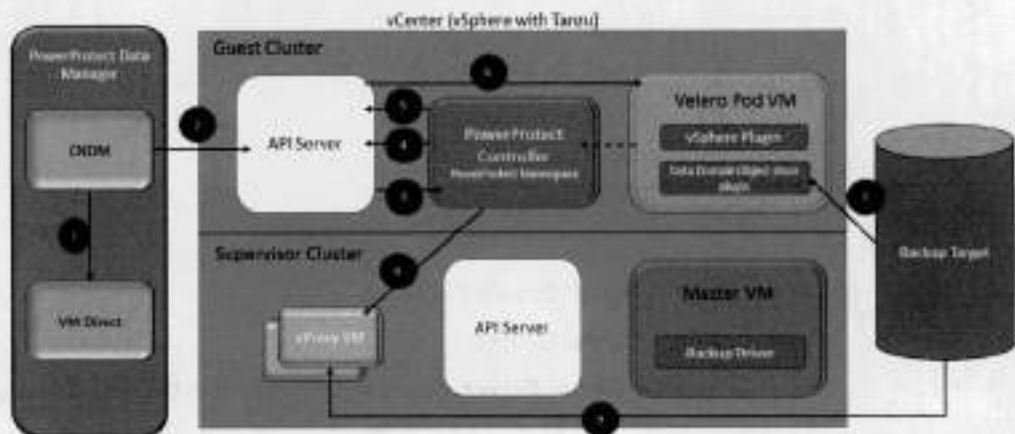
Steps 3, 4, 5: The API server of the guest cluster communicates with the PowerProtect controller (PowerProtect namespace). The BackupJobCR is created and, in turn, the PowerProtect controller sends Velero BackupCR to the API server. The API server then communicates with Velero PodVM (Velero namespace). The API server sends BackupCR to the vSphere plug-in within Velero PodVM.

Steps 6, 7: Communication between the supervisor and guest cluster happens with the help of the vSphere plug-in. The vSphere plug-in within Velero PodVM of the guest cluster initiates communication with the API server of the supervisor cluster by sending Backup CR. Once the supervisor cluster has the information about the guest cluster's Backup CR, the API server communicates to the backup driver, which is a component of the Velero Master VM, in order to take an FCD snapshot. The backup driver is responsible for creating and deleting snapshots that are backed by CSI volumes.

Steps 8, 9, 10: Once the FCD snapshots are taken, the API server of the supervisor cluster communicates back to the vSphere plug-in (Velero PodVM) of the guest cluster that the snapshot has been completed. The PowerProtect controller learns about the FCD snapshot information and communicates with vProxy VM. PowerProtect controller creates a session with vProxy VM, which is created in advance. When the session is established between the controller and vProxy, vProxy learns about the snapshot information from the backup driver and moves the data from FCD to the backup target.

Restore workflow

The restore process for a TKG cluster is manually triggered through the PowerProtect Data Manager UI. The process is similar to the backup process.



Steps 1, 2: CNDM interacts with VM Direct to find and reserve vProxy, which is pre-installed and is responsible for moving data from the backup target to FCD. CNDM then initiates communication with the API server of the guest cluster by sending RestoreJobCR that contains the restore request information. It also passes details to the API server about which vProxy is being used to restore the data.

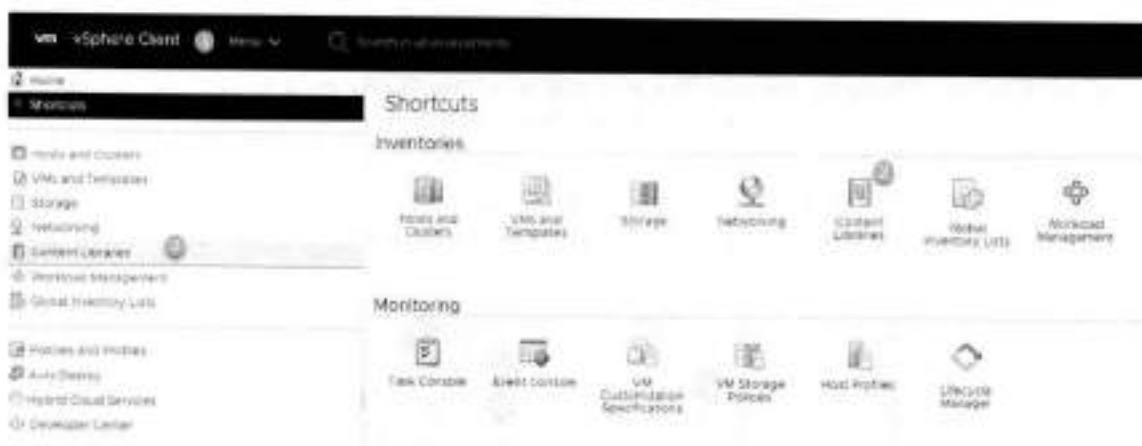
Step 7, 8, 9: Velero PodVM restores the metadata from the backup target and communicates with the PowerProtect controller. The PowerProtect controller asks the vProxy VM to restore the data using the backup driver. Once the restore is completed, the guest cluster is notified about the job completion.

Configuring VMware vSphere with Tanzu

Introduction In order to provision Tanzu Kubernetes cluster, the content library must be created in the vCenter server that manages the vSphere cluster where the supervisor cluster runs.

Create content library The content library provides the distribution of Tanzu Kubernetes releases in the shape of OVA templates.

1. Log in to the **vCenter Server** with administrator credentials.
2. Select **Content Libraries** under **Inventories**.



- **Name:** Specify the name.
 - **Notes:** Optional.
 - **vCenter Server:** Select the vCenter from the drop-down list (if there are multiple vCenterservers).
4. Click **NEXT**.

Note: You can create a Subscribed Content Library to automatically put the latest released images, or you can create a Local Content Library and upload the images manually. Subscription URL: <https://wp-content.vmware.com/v2/latest/lib.json>

5. Click **NEXT**.
6. Verify the identity of the subscription host and click **YES** to proceed.
7. Select the storage location for the library contents and click **NEXT**.
8. Review the content library settings and click **FINISH**.

The library is created and is available under the **Content Libraries** section.

9. Click the library and review the details and OVAs available.



The HA proxy is an appliance that is deployed as an OVA. The HA proxy is used to route network traffic between Kubernetes pods and acts as load balancer as well. To enable workload management, networking must be configured. There are two options—NSX-T Data Center or vSphere Distributed Switch (vDS) networking with external load balancing. Currently, the HA proxy is supported and must be configured before workload management is enabled. Workload management cannot be set up without the HA proxy instance.

For requirements to enable networking, see <https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-B1388E77-2EEC-41E2-8681-5AE549D50C77.html#GUID-B1388E77-2EEC-41E2-8681-5AE549D50C77>.

To deploy the HA proxy and configure IP addresses, ensure that the OVA is available at the local or content library. To find the location from which to download the OVA, see the following VMware knowledge base article: https://kb.vmware.com/s/article/80735?lang=en_US&queryTerm=80735.

1. Log in to **vCenter server** with administrator credentials.
2. Click **Content Libraries** and select the created library (which should have HA proxy OVA).
3. Right-click the OVA, select **New VM from This Template**, and provide the following details:
 - **Select a name and folder:** Specify the VM name and select the location for the VM.
 - **Select a compute resource:** Select the destination compute resource for this operation.
 - Review details and click **NEXT**.
 - Check **I accept all license agreements** and click **NEXT**.

- **Configuration:** Select **Default** or **Frontend Network** as per the network design.

Note: To deploy the appliance with 2 NICs, that is, a management and single workload network, select **Default**; to deploy the appliance with 3 NICs, that is, a management network, a single workload network, and a dedicated front-end network, select **Frontend Network**.

4. Select the storage for the configuration and disk files.
5. Select a destination network for each source network.
6. Customize template: Provide root password, network configuration, and load balancing IP ranges.
7. Verify the configuration and click **FINISH**.

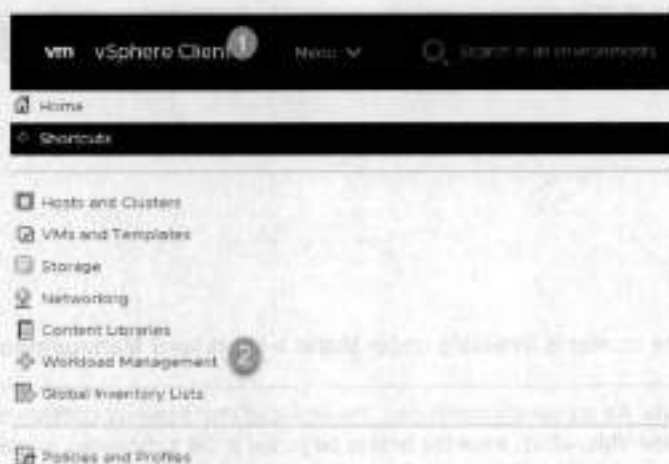
Enable workload management

Enabling workload management on a vSphere cluster creates a supervisor cluster. Workload management enables deploying and managing Kubernetes workloads in vSphere. By using workload management, you can leverage both Kubernetes and vSphere functionality. Once vSphere cluster for workload management is configured, namespaces can be created, which provides compute networking and storage resources for Kubernetes applications.

- **Network Support:** You can select between two networking stacks when configuring workload management such as NSX-T and vCenter server networks. Click **Menu > Workload Management > Network Support**.
- **HA and DRS Support:** HA and DRS must be enabled on the vSphere cluster in fully automated mode on the cluster where you set up workload management.
- **Storage Policy:** Storage policies must be created to specify the datastore placement of the Kubernetes control plane VMs, containers, and images.
- **Load Balancer:** If you use the vCenter Server network, you must configure a load balancer to support the network connectivity to workloads from client networks and to load-balance traffic between Tanzu Kubernetes clusters. The type of load balancer supported is HA Proxy.
- **Tanzu Kubernetes Grid:** The content library must be created on the vCenter server system. The VM image that is used for creating the nodes of Tanzu Kubernetes clusters is pulled from that library. This library will contain the latest distributions for Kubernetes and another OS (<https://wp-content.vmware.com/v2/latest/lib.json>).

To enable workload management:

1. Log in to the vCenter server with administrator credentials.
2. Select **Workload Management**.



3. Click **GET STARTED**.

4. Make the following selections:

- **vCenter Server and Network:** Select a vCenter and then select a networking stack option and click **NEXT**.
- **Select a Cluster:** Select the compatible cluster listed in the cluster details and click **NEXT**.
- **Control Plane Size:** Allocate capacity for the Kubernetes control plane VMs. The amount of resources that you allocate to the control plane VMs determines the amount of Kubernetes workloads that the cluster can support. Select the resource allocation size and click **NEXT**.
- **Storage:** Select the storage policy to be used for datastore placement of Kubernetes control plane VMs and containers. The policy is associated with a datastore on the vSphere environment.
- **Load Balancer:** The load balancer must be configured to support network connectivity to workloads from client networks and to load-balance traffic between Tanzu Kubernetes clusters. The type of load balancer supported is HA proxy.
- **Management Network:** The workload management consists of three Kubernetes control plane VMs and the spherelet process on each host, which allows the host to be joined in a Kubernetes cluster. The cluster where the workload management is connected to management network supports traffic vCenter server.
- **Workload Network:** The network to support traffic to the Kubernetes API and to workload and services.
- **TKG configuration:** Add the content library to give access to the workloads.
- **Review and Confirm:** Review all the details before confirming the setup for workload management on the cluster.

5. Click **FINISH**.



The cluster is available under **Menu > Workload Management > Clusters**.

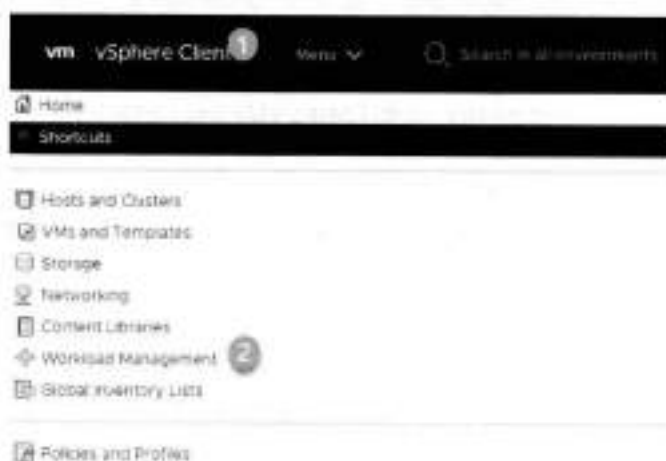
Note: As previously mentioned, the workload management consists of three Kubernetes control plane VMs, which allow the host to be joined in the Kubernetes cluster. Once the workload cluster is created, three SupervisorControlPlaneVMs are created. These VMs are actually Kubernetes nodes that interact with the vSphere infrastructure in order to provide the services and capabilities for vSphere with Tanzu.

Create supervisor namespaces

Once the supervisor cluster is deployed, configured, and licensed, the supervisor namespace can be deployed on the supervisor cluster to run Kubernetes applications.

To create supervisor namespaces:

1. Log in to the vCenter server with administrator credentials.



Workload Management

Namespaces Clusters Updates

ADD CLUSTER DISABLE EXPORT LOGS

Cluster	Namespaces	Pods
Cluster-workload	9	3

3. Select the supervisor cluster.
4. Click **Namespaces**.



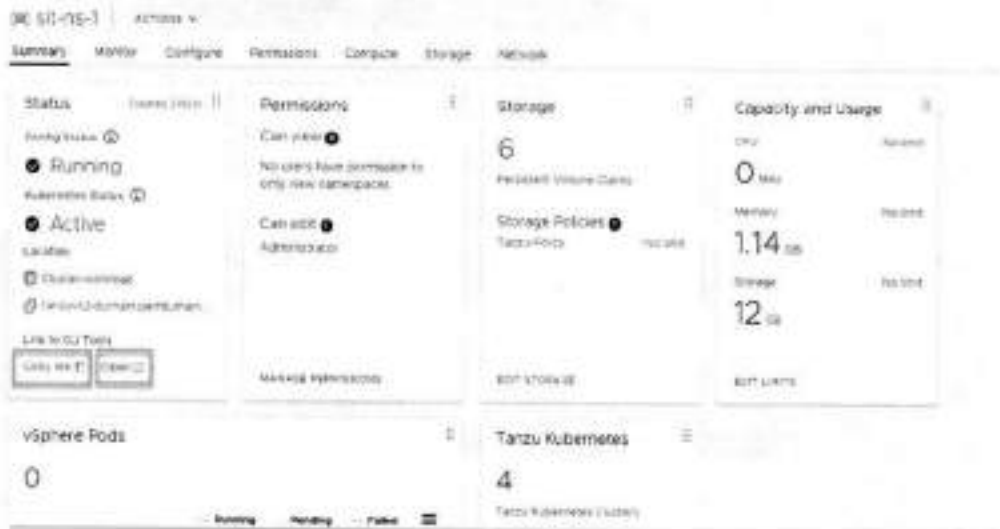
5. Click **NEW NAMESPACE**.
6. Select the cluster where you want to create the namespace.



7. Provide the name.
8. Click **CREATE**.

The namespace has been created and is available under **Menu > Workload Management > Namespaces**.

To access the namespace, you must have the Kubernetes CLI tool installed as a plug-in. You can get that CLI tool by clicking **Copy Link** or **Open**.



Resource Limits and **Object Limits** details are available through vCenter server under **Configure**.



Storage Policies, Config Map, Secrets, and Persistent Volume Claims details are available under the **Storage** section.

The screenshot shows the Tanzu console interface for namespace 'sit-ns-1'. The 'Storage' tab is selected, displaying a table of Storage Policies. A sidebar menu on the left lists 'Storage Policies', 'Config Map', 'Secrets', and 'Persistent Volume Claims'. The table shows one policy named 'Tanzu-Policy' with an available capacity of 14 TB.

Storage Policy	Available Capacity
Tanzu-Policy	14 TB

Network Policies, Services, Ingress and Endpoint information is available under the **Network** section.

9. Download the CLI plug-in.

Kubernetes CLI Tools

Kubectl + vSphere plugin

Download the CLI tools package to view and control namespaces in vSphere. [Learn more](#)

SELECT OPERATING SYSTEM*

DOWNLOAD CLI PLUGIN WINDOWS

Check out CLI plugin Windows



Get started with CLI Plugin for vSphere

Kubernetes CLI tool lets you manage your namespaces. Below are a few steps that will help you get started.

1. Verify that the SHA256 checksum of `vsphere-plugin.zip` matches the checksum in the provided file `sha256sum.txt`. In Powershell run command `Get-FileHash -Algorithm SHA256 -Path vsphere-plugin.zip` to display the checksum
2. Put the contents of the .zip file in your OS's executable search path
3. Run command `kubectl vsphere login --server=IP_or_hostname` to log in to server
4. Run command `kubectl config get-contexts` to view a list of your Namespaces
5. Run command `kubectl config use-context <context>` to choose your default context

You can access the namespaces and create guest clusters using the CLI tool.

Create Tanzu Kubernetes cluster (guest cluster)

Tanzu Kubernetes cluster is created by invoking a Tanzu Kubernetes Grid service declarative API. Once the cluster is created, you can manage and deploy workloads to it by running the Kubectl command.

To create TKG clusters:

1. Download and install the [Kubernetes CLI tool](#) for vSphere, as mentioned in the previous section.
2. Log in to the namespace context as follows:

```
kubectl-vsphere.exe login --insecure-skip-tls-verify --
vsphere-username
```

3. Verify the control plane and storage class:

```
kubectl get
kubectl get sc
```

4. Switch context to the supervisor namespace where you want provision Tanzu Kubernetes Cluster:

```
Kubect config get-contexts
```

5. Construct the YAML file for provisioning Tanzu Kubernetes Cluster and save it as `<cluster-name.yaml>`; for example:

```
apiVersion: run.tanzu.vmware.com/v1kind:
TanzuKubernetesCluster metadata:
name: tkg-cluster-0namespace: test-ns-1
spec:
distribution: version: v1.18.5
topology: controlPlane:
count: 1
```

6. Provision the cluster by running the apply command:

```
Kubectl apply -f <cluster-name>.yaml
```

7. Verify the provisioned cluster:

```
Kubectl cluster-info
Kubectl get nodes
Kubectl get ns
Kubectl api-resources
```

8. At step 5, the yaml file specifies 1 control plane and 3 worker nodes. This can be verified at vCenter under **Namespaces**.



Configuring PowerProtect Data Manager to protect Tanzu Kubernetes workloads

Introduction

PowerProtect Data Manager discovers the Tanzu Kubernetes clusters by using the IP address or Fully Qualified Domain Name (FQDN) of the vCenter server where the Tanzu Kubernetes clusters reside. PowerProtect Data Manager uses the discovery service account and the token (kubeconfig file) to integrate with kube-APIserver of the Kubernetes cluster.

Asset discovery

The Tanzu Kubernetes clusters are hosted by vCenter server; therefore, the vCenter server gets registered with PowerProtect Data Manager. PowerProtect Data Manager discovers the namespaces as assets.

1. Log in to the PowerProtect Data Manager UI with administrator credentials.
2. On the left pane of the PowerProtect Data Manager UI, click **Infrastructure**.



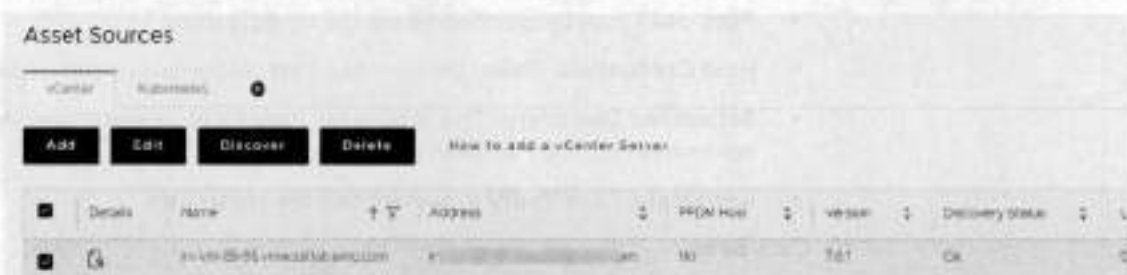
4. Select **vCenter**.
5. Click **Add**.
 - **Name:** Specify the name.
 - **FQDN/IP:** Specify the IP address or FQDN.
 - **Port:** 443
 - **Host Credentials:** Select credentials from the drop down if already added, if not, click **AddCredentials** and provide:
 - FQDN of vCenter server
 - Username
 - Password
 - **vSphere Plugin:** Check box to install vSphere plug-in.
 - **Scheduled Discovery:** This is optional; toggle if you need to specify automated discovery at given time schedule.

The screenshot shows the 'Edit vCenter' configuration window. It includes the following fields and options:

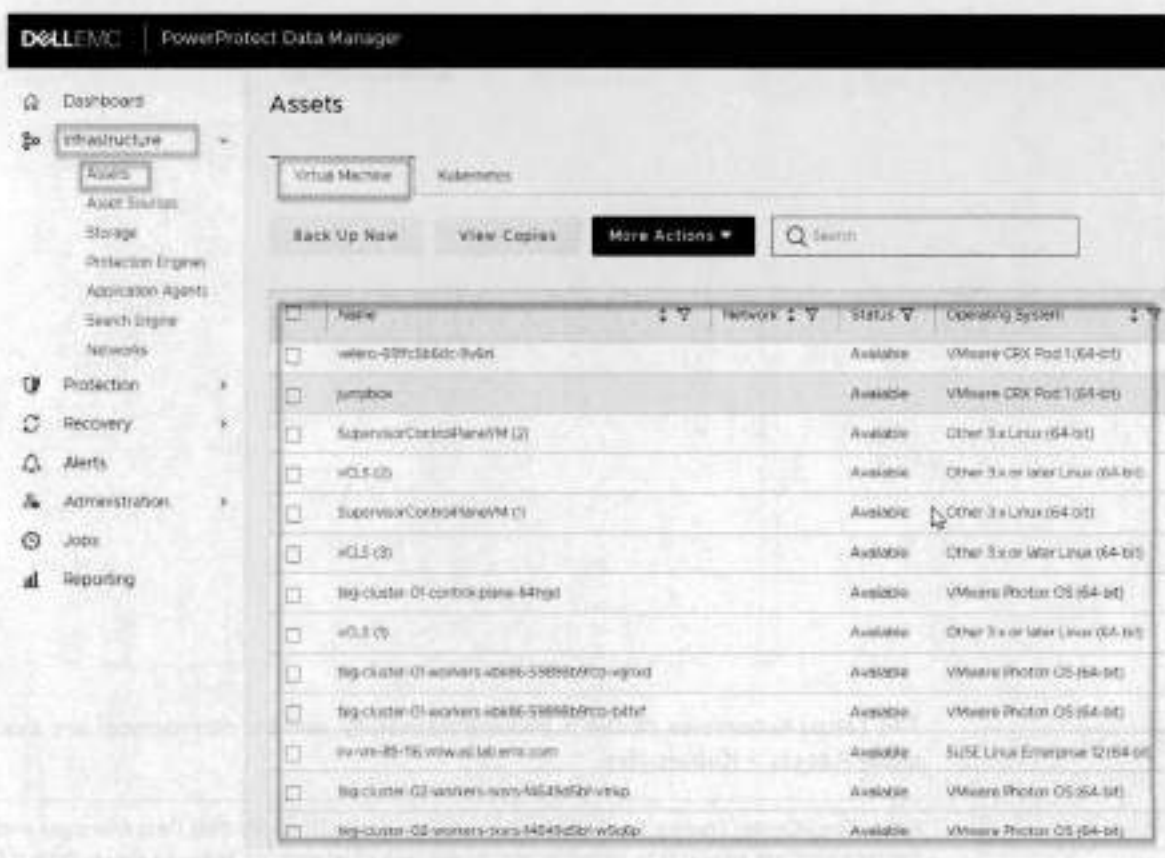
- Name:** A text input field.
- FQDN/IP:** A text input field.
- Port:** A text input field containing the value '443'.
- Host Credentials:** A dropdown menu.
- Schedule Discovery:** A toggle switch that is currently turned off.
- Discovery Time:** Three dropdown menus for selecting the day of the week, month, and AM/PM.
- Certificate:** A section showing a certificate icon and the text 'Verified'.
- Buttons:** 'Cancel' and 'Save' buttons at the bottom right.

6. Click **Save**.

vCenter server is successfully registered with PowerProtect Data Manager and is available under **Infrastructure > Asset Sources**.



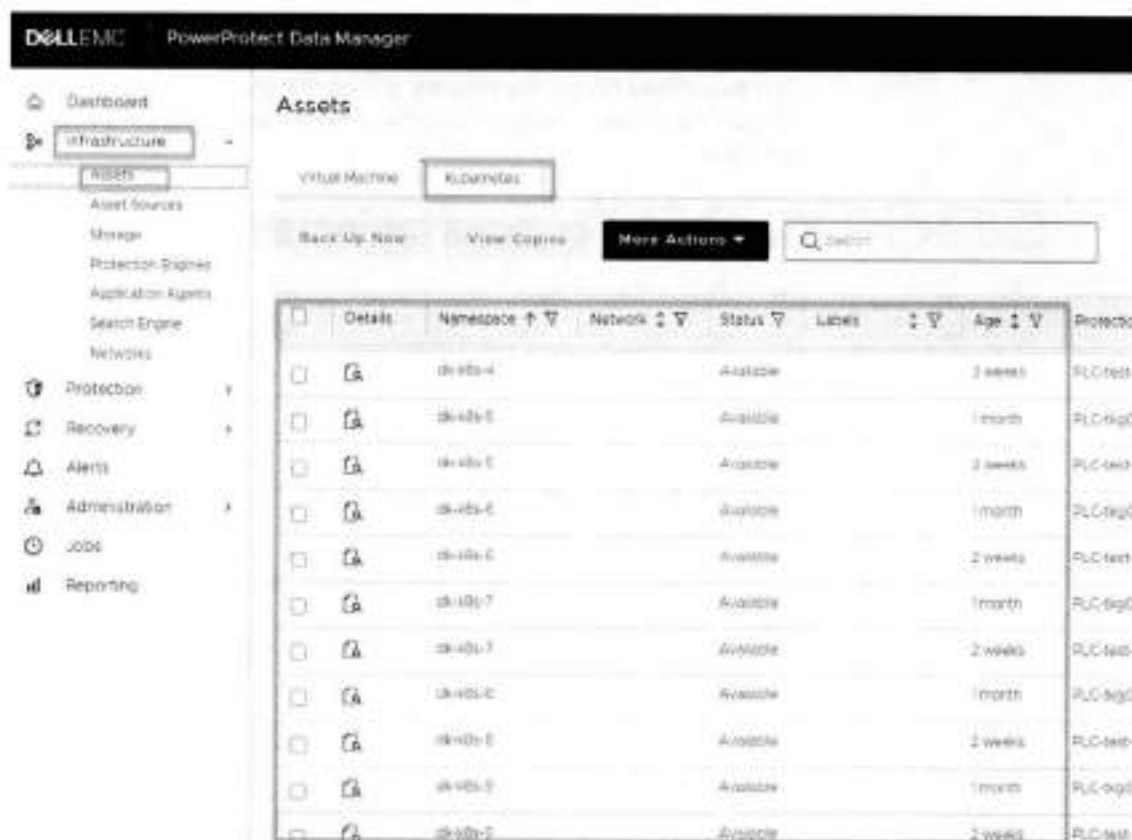
The associated supervisor clusters and control plane VMs with the vCenter are available as assets and are available under **Infrastructure > Assets > Virtual Machine**.



7. To add Tanzu Kubernetes clusters, select **Infrastructure**.
8. Click **Asset Sources** and select **Kubernetes**.
9. Click **Add**.
 - **Tanzu Cluster:** Click the toggle button to enable Tanzu Cluster.
 - **Select vCenter:** Select the vCenter that contains the Tanzu Kubernetes cluster.
 - **Name:** Specify the name.
 - **FQDN/IP:** Provide the IP address or FQDN of the cluster.

- **Port:** 6443 (can be changed as per the configuration).
- **Host Credentials:** Select the specified host credentials or add credentials.
- **Scheduled Discovery:** This is optional; toggle if you need to specify automated discovery at given time schedule.
- **Certificate:** Click **Verify** to authenticate the credentials.

10. Click **Save**.



The Tanzu Kubernetes cluster is added successfully, and the namespaces are available under **Assets > Kubernetes**.

Note: The vCenter hosting Tanzu Cluster is registered with PowerProtect Data Manager and the namespaces are available to be protected. At the time of integration between PowerProtect Data Manager and Tanzu Kubernetes cluster, **velero-ppdm** and **powerprotect**, the two namespaces are created on Tanzu Kubernetes cluster.

kube-public	Active	16d
kube-system	Active	16d
mysql-alter-1	Active	12d
mysql-ns-1	Active	14d
mysql-ns-multi-3	Active	14d
powerprotect	Active	3d3h
recover-dk-k8s-8	Active	11d
tkg01-new	Active	5d9h
velero-ppdm	Active	3d3h
velero-vsphere-plugin-backupdriver	Active	14d
vmware-system-auth	Active	16d
vmware-system-cloud-provider	Active	16d
vmware-system-csi	Active	16d

The **powerprotect** namespace contains powerprotect-controller pod, and **velero-ppdm** namespace contains velero pod and backup-driver pod. The backup-driver pod is responsible for taking and deleting snapshots for paravirtualized CSI volumes of Tanzu Kubernetes clusters.

Enable Velero vSphere Operator

The Velero supervisor namespace contains a Velero pod that helps to back up metadata and stage it to the target storage in the case of a bare metal environment. It also performs PVC and metadata backup in the case of VMware Cloud Native Storage (vCNS). Velero vSphere Operator must be configured and enabled to perform backups. PowerProtect Data Manager uses the Velero plug-in to protect the Kubernetes workloads. PowerProtect Data Manager supports existing Velero deployment in the default namespace and also accesses directly Velero auto-deployed in the **velero-ppdm** namespace by PowerProtect Data Manager.

1. Verify auto-deployed namespaces by PowerProtect Data Manager:

```
kubectl get pods -n powerprotect
```

```

root@vcenter-12-100 -11 [kubecli get pods -n powerprotect]
NAME                                READY   STATUS    RESTARTS   AGE
powerprotect-controller-85175fd805-wjwz  1/1     Running   0           3d3h
root@vcenter-12-100 -11 [kubecli get pods -n velero-ppdm]
NAME                                READY   STATUS    RESTARTS   AGE
backup-driver-7d8b00fe995-xf0ms        1/1     Running   0           3d3h
velero-9ddb577f9-4ncqg                 1/1     Running   0           3d3h

```

2. Log-in to the vCenter server with administrator credentials.
3. Select the workload cluster.
4. Click **Configure**.
5. Expand **Supervisor Services**.
6. Click **Services**.
7. Select **Velero vSphere Operator**.
8. Click **Enable**.



9. Verify that the Velero vSphere Operator service is enabled.

VM direct engine

VM direct is the protection engine within PowerProtect Data Manager. It is used to manage and protect VM assets. VM direct engine are created manually:

1. Log in to the PowerProtect Data Manager UI with administrator credentials.
2. On the left pane of the PowerProtect Data Manager UI, click **Infrastructure**.
3. Click **Protection Engines**.
4. Click **Add** to add new protection engine.
5. Specify the details for the protection engine:
 - **Hostname:** Provide the hostname.
 - **Gateway:** Specify the gateway.
 - **IP address:** Provide the valid IPv4 or IPv6 address.
 - **Netmask:** Provide the valid netmask address.
 - **vCenter to Deploy:** Select the vCenter on which the VM is deployed.
 - **ESX Host/ Cluster:** Select the ESX host or cluster on which the VM is deployed.
 - **Supported Protection Type:** Click the drop-down and select **Kubernetes**.
 - **Primary DNS:** Specify the primary DNS.
 - **Secondary DNS:** Specify the secondary DNS (if any).
 - **Tertiary DNS:** Specify the third DNS (if any).
 - **Network:** Click the drop-down and select the required VM network.
 - **Data Store:** Select the datastore.
 - **Transport Mode:** Select the required transport mode—hot add, Network Block or Hot add, Network Block Device.
6. Specify the network configuration details.
7. Verify the summary and click **Save**.

The available VM direct engines are shown under the **VM Direct Protection Engines** tab.

Backup configuration

Asset backups can be taken manually or scheduled. A protection policy must be created first:

1. Log in to PowerProtect Data Manager UI with administrator credentials.
2. On the left pane of the PowerProtect Data Manager UI, click the **Protection** drop-down.
3. Click **Protection Policies**.
4. Click **Add** to add a policy.
 - **Type:** Specify the necessary details:
 - a. Name:
 - b. Description:
 - c. Type: Select **Kubernetes**
 - d. Click **Next**
 - **Purpose:** Select one of the following options, according to the purpose of the backup:
 - a. **Crash Consistent:** Select this option to snapshot persistent volumes bound to the persistent volume claims in the namespace and back them up to the storage target
 - Or**
 - b. **Exclusion:** Select this option to exclude in this group from protection activities and protection rule assistant.
 - c. Click **Next**.
 - **Assets**
 - a. Select the asset(s) to be backed up from the list, or search for a particular asset by typing in the **Find More Assets** field.
 - b. Click **Next**.
 - **Schedule:** In this section, a backup schedule is created, replicated, edited, and deleted.
 - a. Click **+Backup**.
 - b. Fill the required details under the **Add Primary Backup** section.
 - Select **Recurrence** as **Hourly**, **Daily**, **Weekly**, or **Monthly**, and specify the details accordingly.
 - Create Every:
 - Keep For:
 - Start time:
 - End Time:
 - c. Click **Ok**.
 - d. An SLA can be added as follows:

- i Click the drop-down under **SLA**.
 - ii Click **Add** and provide the name and specify the objective(s):
 - **SLA name:**
 - Check box to specify **Recovery Point Option** (minutes, hours, days, weeks, months, or years).
 - Check box if any **Compliance Window** is to be mentioned.
 - Check box to **Verify expired copies are deleted**.
 - Check box to specify **Retention Time Objective** (days, weeks, months or years).
 - Check box to **Verify Retention Lock is enabled for all copies**.
 - iii Click **Save**.
5. Click **Next**.
6. Verify the provided information is correct under the **Summary** section. If yes, click **Finish**.
7. Click **Go to Jobs** to check the status of the policy.

The protection policy triggers the backup at the scheduled time. After a protection policy is created, you can modify the policy with these options:

1. **Edit:** To edit the information or to change the schedule.
2. **Disable:** Backup Schedule is disabled with this option so backup would not be taken.
3. **Export:** Downloadable file that contains the information about the asset protection.
4. **Protect Now:** This option allows you to take backups manually on an ad-hoc basis.
 - **Asset Selection:** Provides options to select the assets:
 - a. All assets defined in the protection policy.
 - b. **Choose some of the assets defined in the policy:** This option allows you to select namespaces within the cluster.
 - **Configuration:**
 - a. This allows you to select the type of backup:
 - **Full:** Backs up the namespace metadata and persistent volumes and creates a new full backup.
 - **Synthetic full:** Backs up namespace metadata, change blocks for persistent volumes on VMware first class disks, all other persistent volumes and creates a new full backup.
 - b. **Keep For:** In days
 - c. Click **Next**.
 - **Summary:** Verify the information.
5. Click **Backup**.

Monitor the backup job by clicking **Go to Jobs**.



The backup job is completed successfully, and the details such as task ID, vProxy, throughput, storage, and so on are available in the backup job details.

Replication configuration

The replication is configured with the existing backup policy or a new policy.

To configure replication on the existing backup policy:

1. Log in to the PowerProtect Data Manager UI with administrator credentials.
2. On the left pane of the PowerProtect Data Manager UI, click the **Protection** dropdown.



4. Select the existing backup policy.
5. Click **Edit** and click the **Back** button.
6. Select the **Primary Backup** schedule.

Configuring PowerProtect Data Manager to protect Tanzu Kubernetes workloads



7. Click **Replicate**.
8. Add Primary Replication details:
 - **Replicate Every**: Provide the appropriate time.
 - **Keep for**: Specify the number of days.
 - **Start Time**
 - **End Time**



9. Click **OK**.
10. Verify that the replicate schedule is created.

11. Click **Finish**.
12. Verify that the replication job is created under the **Jobs** section.
13. To run the replication now, select existing policies from **Protection > Protection Policies**.



14. Click **Protect Now**.
 - **Asset Selection:** Choose one option for ad-hoc protection.
 - **Configuration:** Select the **Replicate Now** option and check box to select replication storage.
 - **Summary:** Click **Protect Now** to start replication.
15. Click **Go to Jobs** to view the progress and details.



16. Verify that the replication job is successfully completed. Click the details button for detailed results.

Restore configuration

The recovery of the assets is a manual process. The restoration of the Kubernetes namespaces, which includes pods, stateful sets, PVCs, and other resources, is done at the same cluster. PowerProtect Data Manager provides options to recover the Kubernetes namespaces at the same cluster as well as the alternate cluster.

1. Log in to the PowerProtect Data Manager UI with administrator credentials.
2. On the left pane of the PowerProtect Data Manager UI, click **Recovery**.
3. Click **Assets**.
4. Click **Kubernetes** on top and select the namespaces to be restored.
5. Click **Restore**.
 - **Select Copy:**
 - a. Select the restore copy. The most recent copy is used by default. To change from the default copy, click **Change Copy**.

- b. Click **Ok**.
- c. Click **Next**.
- **Cluster:** This provides the option to select the cluster on which assets are to be restored.
 - a. **Restore to Original Cluster:** The assets are restored to the source cluster from which the backup is taken.
 - b. **Restore to Alternate Cluster:** The assets are restored on the alternate cluster. To use this option, the alternate cluster is added as an asset source to the managing PowerProtect Data Manager.
- **Purpose:** Select the option for what is to be restored.
 - a. **Restore Namespace and Select PVCs:** This option restores the namespace and a subset of PVCs in the namespace. The namespace resources, including pods, services, secrets, and deployments, will not be overwritten during a restore. All other resources that do not currently exist in the namespace will be restored.
 - b. **Restore PVCs only:** This option will restore only PVCs.
- **Restore Type:** Restore type has different options depending on the purpose of the restore.
 - a. If the purpose is to restore namespaces and PVCs, the options are:
 - Restore to Original Namespace
 - Restore to New Namespace
 - Restore to an Existing Namespace
 - b. If the purpose is to restore PVCs only, the options are:
 - Restore to Original Namespace
 - Restore to an Existing Namespace
- **PVCs:** Select PVCs to be restored to the namespace. The options are:
 - a. Overwrite content of existing PVCs.
 - b. Skip restores of existing PVCs.
- **Summary:** Verify the information and click **Restore**.

Note: When the restore job starts, a new namespace is created at the Kubernetes cluster. Once the restore job is complete, the pod and PVC get created with details of pod and PersistentVolume bound to PVC.

Technical support and resources

Technical support

[Dell.com/support](https://www.dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage technical documents and videos](#) provide expertise that helps to ensure customer success on Dell Technologies storage platforms.

Related resources

For more information, see the following resources:

- vSphere with Tanzu brief: <https://d1ftc35gcffzn.cloudfront.net/tanzu/VMware-Tanzu-Basic-Solution-Brief.pdf>
- vSphere with Tanzu configuration: <https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-152BE7D2-E227-4DAA-B527-557B564D9718.html>
- vSphere distributed switch with HA proxy <https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-C86B9028-2701-40FE-BA05-519486E010F4.html>
- Storage policy creation with vSphere : <https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-544286A2-A403-4CA5-9C73-8EFF261545E7.html#GUID-544286A2-A403-4CA5-9C73-8EFF261545E7>
- Kubernetes CLI tool <https://docs.vmware.com/en/VMware-vSphere/7.0/vmware-vsphere-with-tanzu/GUID-0F6E45C4-3CB1-4562-9370-686668519FCA.html>
- Dell EMC PowerProtect Data Manager Data Sheet: <https://www.dell.com/en-me/collaterals/unauth/data-sheets/products/data-protection/h17691-dell-emc-powerprotect-software-ds.pdf>
- Dell EMC PowerProtect Data Manager: <https://www.delltechnologies.com/en-in/data-protection/powerprotect-data-manager.htm#scroll=off>
- CSI-driver-host-path: <https://github.com/kubernetes-csi/csi-driver-host-path>

Multi-Cloud Data Services for Dell EMC PowerProtect

Abstract

This document explains how a data protection offering from Dell Technologies™ and Faction provides a fully managed service to protect on-premises and multi-cloud environments.

April 2021

Revisions

Revisions

Date	Description
April 2021	Initial release

Acknowledgments

Authors: Vinod Kumaresan and Parimala Guruprasad

The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

This document may contain certain words that are not consistent with Dell's current language guidelines. Dell plans to update the document over subsequent future releases to revise these words accordingly.

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

Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [4/30/2021] [Technical White Paper] [H18759]

Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents.....	3
Executive summary.....	4
Audience.....	4
1 Introduction.....	5
1.1 Solution benefits.....	6
1.1.1 PowerProtect DD series appliances.....	7
1.1.2 Faction Cloud.....	8
2 Use cases.....	10
2.1 Public cloud protection (backup target).....	11
2.2 Replication target for on-premises DD series and DDVE in the cloud.....	13
2.3 Cyber recovery.....	14
2.3.1 Replicate from DDVE on public cloud to Cyber Recovery vault in the Faction data center.....	15
2.3.2 Replicate from on-premises DD series to Cyber Recovery vault in Faction data centers.....	15
2.3.3 Replicate within Faction data centers.....	16
2.3.4 Compute resource at Faction for Cyber Recovery:.....	16
3 Client connectivity.....	17
4 Example scenarios.....	18
4.1 DD series appliance at Faction as the primary backup target for cloud workloads.....	18
4.2 DD series appliance at Faction as the replication backup target for DDVE deployed on public cloud.....	20
4.3 Mount Faction DD series storage unit on VMs running in multiple cloud environments using BoostFS.....	22
5 Onboarding process and support.....	25
5.1 Implementation.....	25
5.2 Support.....	25
A Technical support and resources.....	26
A.1 Related resources.....	26

Executive summary

As public cloud services continue to grow, the competition between cloud providers drives innovation. As native cloud services evolve, they provide increasingly differentiated value propositions to organizations. A multi-cloud strategy allows users to select the cloud services that best meet their needs, unleashing competitive advantages and productivity gains that would be unattainable with a single cloud. However, with multi-cloud environments, organizations require an integrated data-protection strategy and seamless data movement across their various cloud ecosystems.

Customers are looking for cost efficiency and operational flexibility of their cloud ecosystem without compromising security, data protection, and data integrity that on-premises environments offer.

Multi-Cloud Data Services for Dell EMC™ PowerProtect is a fully managed data-protection-as-a-service solution. It enables customers to back up their workloads across public clouds to a PowerProtect DD series appliance that is hosted in a Faction data center.

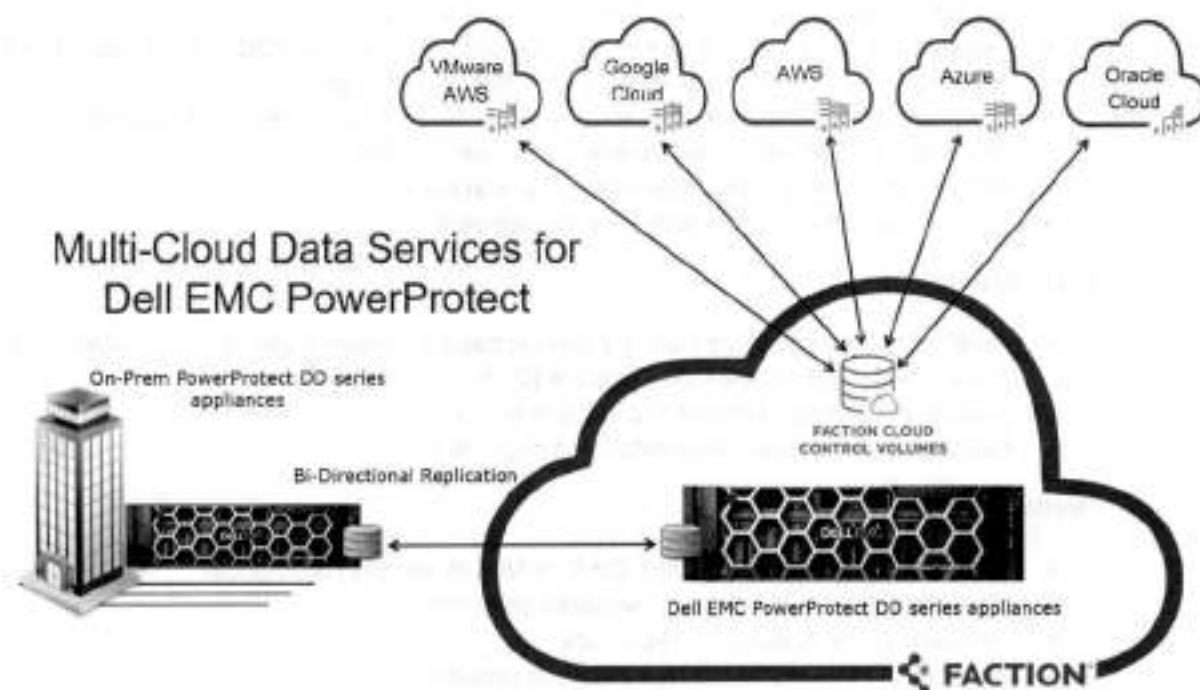
Multi-Cloud Data Services for Dell EMC PowerProtect also integrates with Dell EMC PowerProtect Cyber Recovery, enabling customers to protect their organization from ransomware, insider attacks, and other cyber threats. This solution provides an air-gapped and secure Cyber Recovery vault that is hosted in a Faction data center, providing physical isolation of critical customer data.

Audience

This white paper is intended for Dell Technologies customers, partners, and employees looking to protect their on-premises and public cloud workloads in a fully managed, nonpublic cloud destination.

1 Introduction

Multi-Cloud Data Services for Dell EMC PowerProtect is offered through Faction Cloud Control Volumes (CCV) and is backed by Dell EMC PowerProtect DD series appliances. The Multi-Cloud Data Services for Dell EMC PowerProtect solution provides centralized data protection for workloads across various public clouds and on-premises workloads.



Multi-Cloud Data Services for Dell EMC PowerProtect is ideal for the following use cases:

- Backup target for public cloud workloads
- Replication target for on-premises PowerProtect DD series appliances or PowerProtect DD Virtual Edition (DDVE) deployed on the cloud
- Air-gapped, physically isolated vault environment for PowerProtect Cyber Recovery

Multi-Cloud Data Services for Dell EMC PowerProtect offers a low-latency connection (< 2 ms) to the major hyperscale cloud providers. This ability enables efficient backup, archiving, and disaster recovery of both cloud-based and on-premises data workloads.

1.1 Solution benefits

Multi-Cloud Data Services for Dell EMC PowerProtect offers the following benefits to enterprises that require protection for their hybrid-cloud or multi-cloud environments.

Simplified management:

- Enables managing multiple cloud backups in one location
- Removes the requirement to deploy and maintain a PowerProtect DD Virtual Edition (DDVE) appliance in each cloud if it is used as a primary backup target
- Eliminates the maintenance of a secondary data center for replication and Cyber Recovery vault
- Has a single IP, VLAN, or namespace across public clouds
- Deduplicates one copy of data globally across all clouds
- Provides governance and compliance management

Cost savings:

- Up to 75% lower egress cost with Faction contracts for Amazon Web Services (AWS), VMware® Cloud (VMC), and Google Cloud Platform (GCP)
- Zero egress cost for Microsoft® Azure® and Oracle®
- Industry-leading deduplication reduces storage cost¹

Flexibility:

- Native read/write using DD Boost, CIFS, or NFS for various public clouds
- No limitations on quota for cloud-workload protection
- Application mobility across public clouds
- Ability to instantly restore to any multi-cloud provider
- Flexibility for customers to choose technologies from any cloud and not get locked in to one cloud provider

Trusted service:

- Certified Dell Cloud Service Provider
- High-speed access with low-latency connection to all major public clouds

¹ <https://www.delltechnologies.com/en-us/data-protection/powerprotect-backup-appliances.htm>

1.1.1 PowerProtect DD series appliances

The Dell EMC PowerProtect DD series offers the ultimate protection storage appliances that are the latest generation of Dell EMC Data Domain appliances. DD series delivers a fast, secure, and an efficient solution that is optimized for multi-cloud data protection and future demands.

The DD Operating System (DDOS) is the intelligence that powers DD series. It provides the agility, security, and reliability that enables DD series to deliver high-speed, scalable, and industry-leading multi-cloud protection storage for backup, archive, and disaster recovery².

DDOS software elements include the following:

- DD Boost
- DD VTL
- DD Replicator
- DD Cloud Tier
- DD Retention Lock
- DD Encryption
- DD Secure Multi-Tenancy
- DD High Availability
- DD System Manager
- DD Management Center

Dell EMC PowerProtect DD series appliances



DD series can scale up to a physical capacity of 1.5 PB in a single rack, using minimal floor space and lowering power and cooling by up to 41%². DD series provides up to an additional 2 PB of cloud capacity for long-term retention with Dell EMC Cloud Tier.

DD series consists of the DD9900, DD9400, DD6900, DD3300, and a software-defined appliance with PowerProtect DD Virtual Edition (DDVE).

DDVE uses the power of DDOS to deliver software-defined protection storage on-premises and in-cloud. DDVE is fast and simple to download, deploy, and configure, and can be up and running in minutes.

² <https://www.delltechnologies.com/en-in/collaterals/unauth/data-sheets/products/data-protection/h17926-dellemc-powerprotect-dd-ds.pdf>

You can deploy DDVE on any standard hardware, converged or hyperconverged, and run it in VMware vSphere®, Microsoft Hyper-V®, KVM. You can also run DDVE in the cloud with AWS, AWS GovCloud, VMware Cloud, Azure, Azure Government Cloud, and Google Cloud.

A single DDVE instance can scale up to 256 TB in-cloud (AWS, Azure, and Google Cloud) and up to 96 TB on-premises.

As part of Multi-Cloud Data Services for Dell EMC PowerProtect, customers can choose any of the PowerProtect DD series appliance models: DD6900, DD9400, and DD9900 with incremental scaling.

1.1.2 Faction Cloud

Faction, Inc. is a Dell Technologies Platinum Partner and Extended Technologies Complete partner, founded in 2006, and headquartered in Denver, Colorado.

Faction is a leading multi-cloud data-services provider. Faction pioneered cloud-adjacent storage that is powered by patented technology providing data access over low-latency, high-throughput connections to all the major clouds, including AWS, Azure, and Google Cloud Platform.

Faction is also a leading managed service provider for VMware Cloud on AWS, including disaster recovery and production deployments. Also, Faction was the first to offer cloud-attached storage solutions that integrate natively with VMware Cloud on AWS.

Faction Cloud Control Volumes (CCVs) are persistent cloud-attached storage volumes that can connect to public clouds through proprietary, patented network connectivity with ultralow latency. This technology enables organizations to minimize the high egress fees charged by cloud providers. Cloud-attached storage also enables seamless data portability between clouds and avoids vendor lock-in.

Faction services and data centers undergo annual Type II SOC1 and SOC2 and HIPAA compliance audits, with independent outside auditor attestations available under NDA. Faction can create BAA agreements with customers-subject to HIPAA.

The Faction difference:

- **Faction Internetwork Exchange (FIX):** A single IP and single namespace across all clouds and on-premises locations.
- **Application mobility across clouds:** Faction's private and multi-cloud platforms give clients the ability to move, access, scale, and protect data between clouds, without the fear of cloud lock-in. Faction has data centers strategically located next to the public cloud providers (Azure, AWS, VMC, GCP, and OCI).
- **Balancing of cloud network costs:** Faction efficiently blends cloud-connect ports to offer a balance of high and low speeds, driving down TCO. Faction has negotiated contracts with each cloud vendor: \$0 egress charges for Azure and Oracle Cloud.
- **Low-latency connections:** Managed data center locations to offer the lowest latency connectivity to multiple clouds. Faction has dark fiber connections to each of these data centers for low latency and fast connectivity.

Faction provides a portal that allows customers to manage connections to their DD series appliances that are hosted at the Faction data centers. Through the Faction Portal, customers can also provision new connections to cloud providers through the Faction Internetwork Exchange (FIX).

Faction operates seven cloud-service locations (see Figure 1) across the United States, in London, and in Frankfurt.

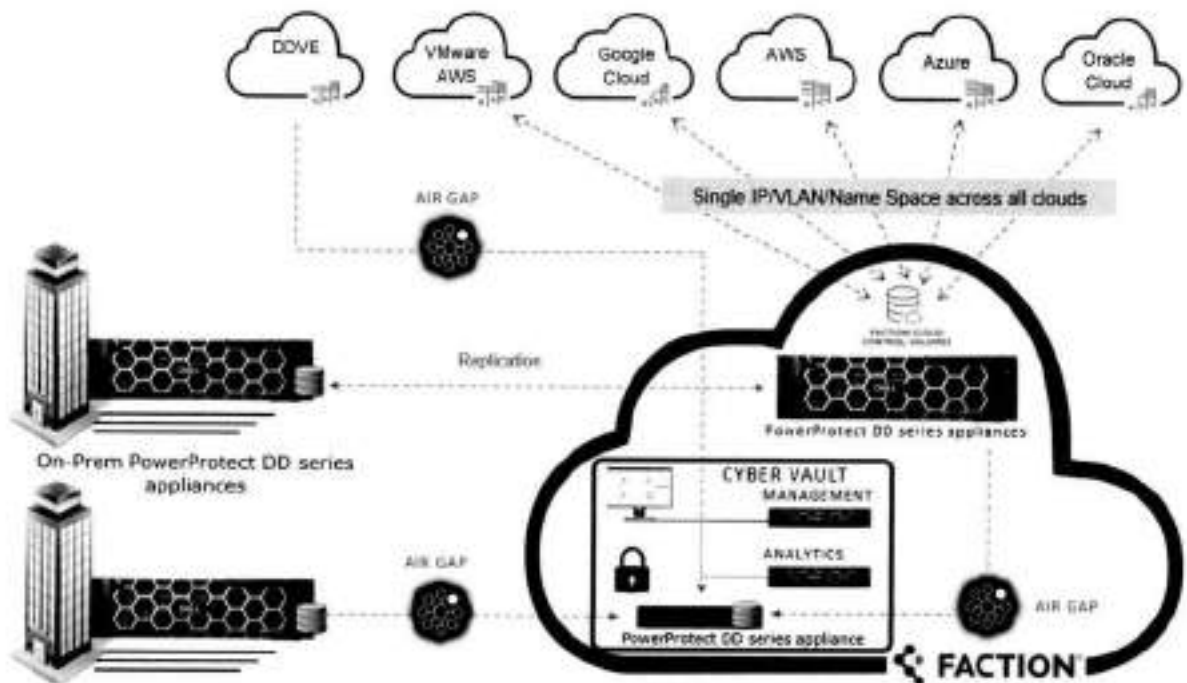


Figure 1 Faction Cloud locations

2 Use cases

Multi-Cloud Data Services for Dell EMC PowerProtect supports the following use cases:

- **Backup target:** Protection of public cloud data to DD series at Faction data centers with a single namespace across multiple clouds.
- **Replication target:** Replicate data from on-premises DD series or DDVE deployed in the cloud for long-term retention (LTR) to Faction data centers or to public clouds, and for Cloud Disaster Recovery.
- **PowerProtect Cyber Recovery:** Enables customers to replicate both on-premises and off-premises workloads to a physically isolated, air-gapped vault.



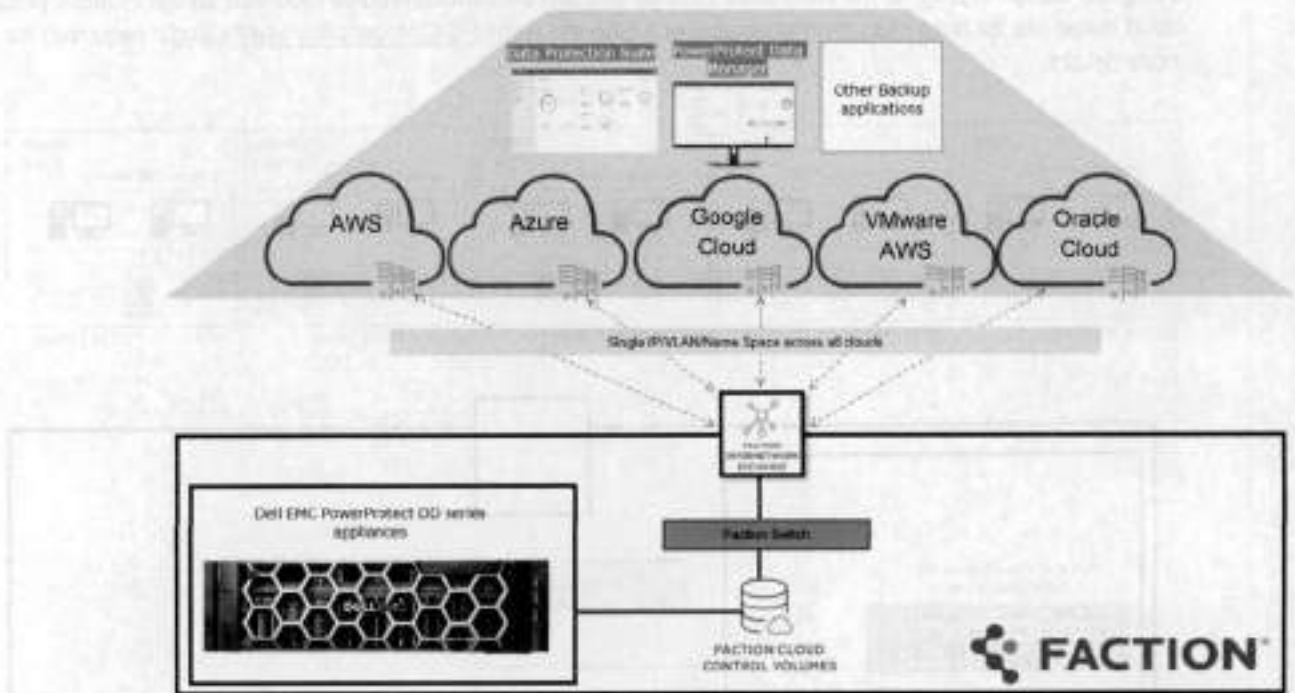
As part of the offer, customers have the option of long-term data retention in Faction's data center and can migrate or replicate data into public clouds. Multi-Cloud Data Services for Dell EMC PowerProtect makes disaster recovery possible using multi-cloud attached storage and compute that is ready in Faction or in supported the public cloud.

Cloud Disaster Recovery (CDR) allows enterprises to copy backed-up VMs from their on-premises environments to the public cloud (AWS and Azure) for the orchestration and automation of DR testing, DR failover and failback of Tier 2 workloads to or from the cloud in a disaster scenario.


Long-Term Retention (Cloud Tier) is available for data on-premises, in the cloud, or both. Multi-Cloud Data Services for Dell EMC PowerProtect can be a remote site to protect data that must be retained for regulatory requirements (governance and compliance) and for workload migration.

2.1 Public cloud protection (backup target)

A DD series appliance at a Faction data center can be used as single backup target to protect workloads across multiple clouds. Using DD Boost, CIFS, and NFS protocols, backup applications can send data to a single DD series appliance at the Faction data center.



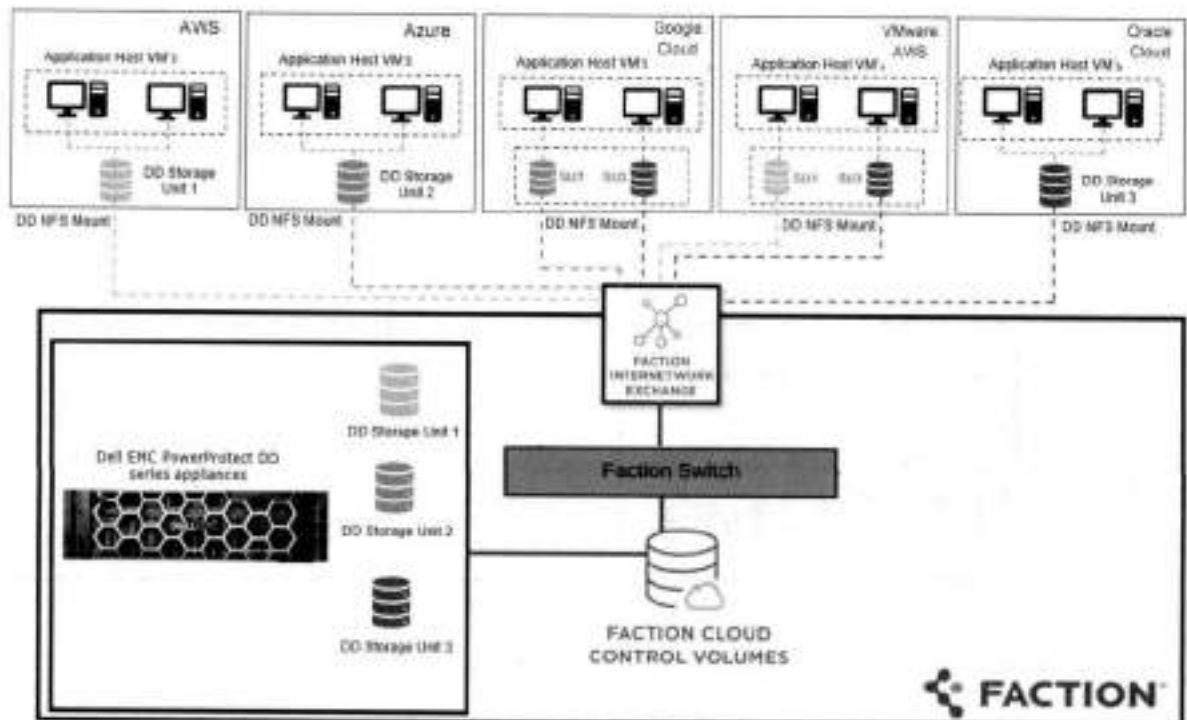
Customers can also use the optional secondary Faction DD series appliance for site replication.



Use cases

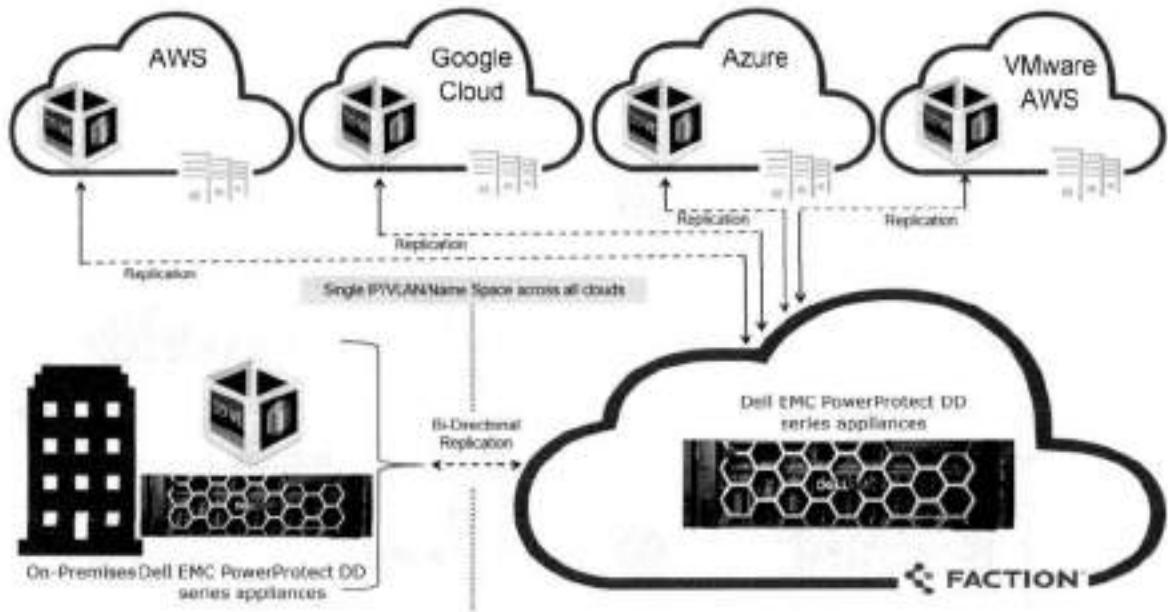
Databases and applications running on public clouds can save backups to the DD series appliances at Faction data centers by mounting the DD series storage unit on the cloud instances using BoostFS plug-in. With direct access to a BoostFS mount point, the application can use source-side deduplication, storage, and network efficiencies of the DD Boost protocol for backup and recovery.

Using DD BoostFS plug-in, the DD series storage unit can be simultaneously mounted on the multiple public cloud instances for data protection and recovery. See the [BoostFS Compatibility Matrix](#) (login required) for more details.



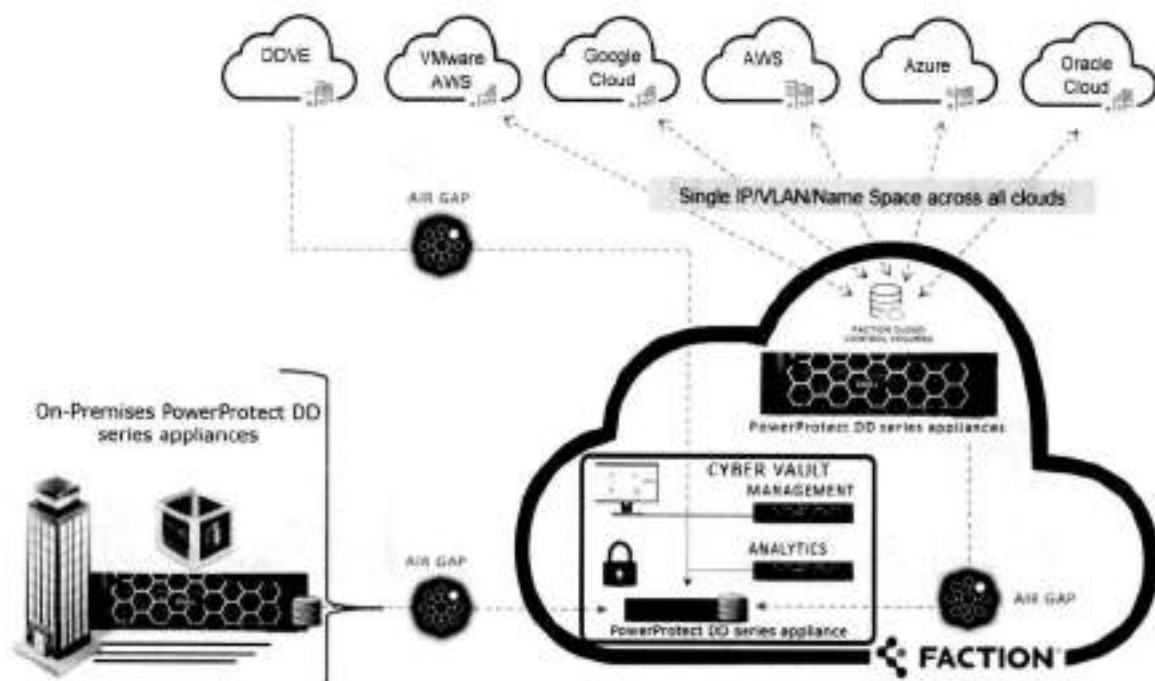
2.2 Replication target for on-premises DD series and DDVE in the cloud

Customers can replicate data from an on-premises DD series or a DDVE in the cloud to a single DD series provisioned in the Faction data centers. This solution eliminates the need for customers to maintain a secondary site for replication. This solution enables long-term retention and disaster recovery.



2.3 Cyber recovery

This secure data-vaulting service is physically and logically air-gapped and is built on secure, multi-cloud-enabled infrastructure that safeguards critical data from cyberattacks.



Combined with the physical security and isolation of the vault, this solution includes an operational air gap. This air gap enables access to the vault only during replication.

At all other times, the vault is disconnected from the client's production environment. Immutable copies of user-selected data are created in the Cyber Recovery vault hosted in the Faction data center. Once a copy of the selected data is safely in the secure, isolated vault, the data cannot be altered, deleted, or otherwise changed for a prescribed duration.

PowerProtect Cyber Recovery is the first solution to fully integrate CyberSense, which adds an intelligent layer of protection to help find data corruption when an attack penetrates the data center. This innovative approach provides full content indexing and uses machine learning to analyze over 100 content-based statistics and detect signs of corruption due to ransomware. CyberSense finds corruption with up to 99.5% confidence, helping you identify threats and diagnose attack vectors while protecting your business-critical content—all in the security of the vault.

When data recovery is required, the data from the Cyber Recovery vault can be restored to AWS, VMC on AWS, Microsoft Azure, Google Cloud, Oracle Cloud, or back to the on-premises environment.


Data in DD series appliances can be replicated to a Cyber Recovery vault in the Faction data center from the following:

- Public cloud
- On-premises
- Within the Faction cloud

Use cases


2.3.1 Replicate from DDVE on public cloud to Cyber Recovery vault in the Faction data center

For cloud-native applications using DDVE, the Cyber Recovery vault service enables customers to replicate critical data to a secure Cyber Recovery vault at a Faction data center.



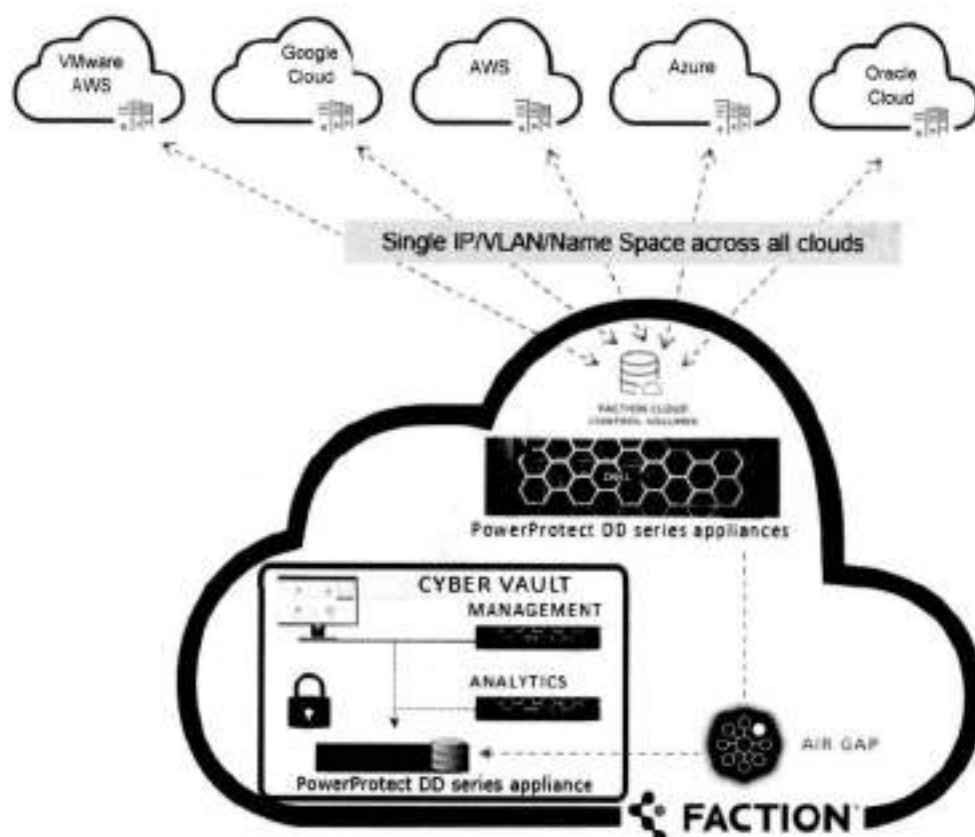
2.3.2 Replicate from on-premises DD series to Cyber Recovery vault in Faction data centers

Customers can replicate data from an on-premises DD series appliance to a Cyber Recovery vault in the Faction data center.



2.3.3 Replicate within Faction data centers

Data residing in any public cloud can be backed up to the DD series appliance in a Faction data center as a primary target. These backups can be replicated to a secure Cyber Recovery vault in a Faction data center.



2.3.4 Compute resource at Faction for Cyber Recovery:

Compute resources required for Cyber Recovery are deployed on a dedicated VMware ESXi™ host or hosts. This VMware compute environment includes the following:

- Required ESXi host to support management VMs
 - Available resources: 50 GHz, 250 GB RAM, 22 TB storage
- Optional ESXi CyberSense analytics host or hosts
 - Available resources: 95 GHz, 635 GB RAM, 14 TB storage

The number of analytics hosts required varies based on the amount of data to be analyzed.

All VMware licensing is included in the cost of these hosts. Customers are provided VMware vCenter® access to install and manage all VMs that run in the VMware compute environment.

Cyber Recovery and the optional CyberSense analytics licenses must be purchased separately.

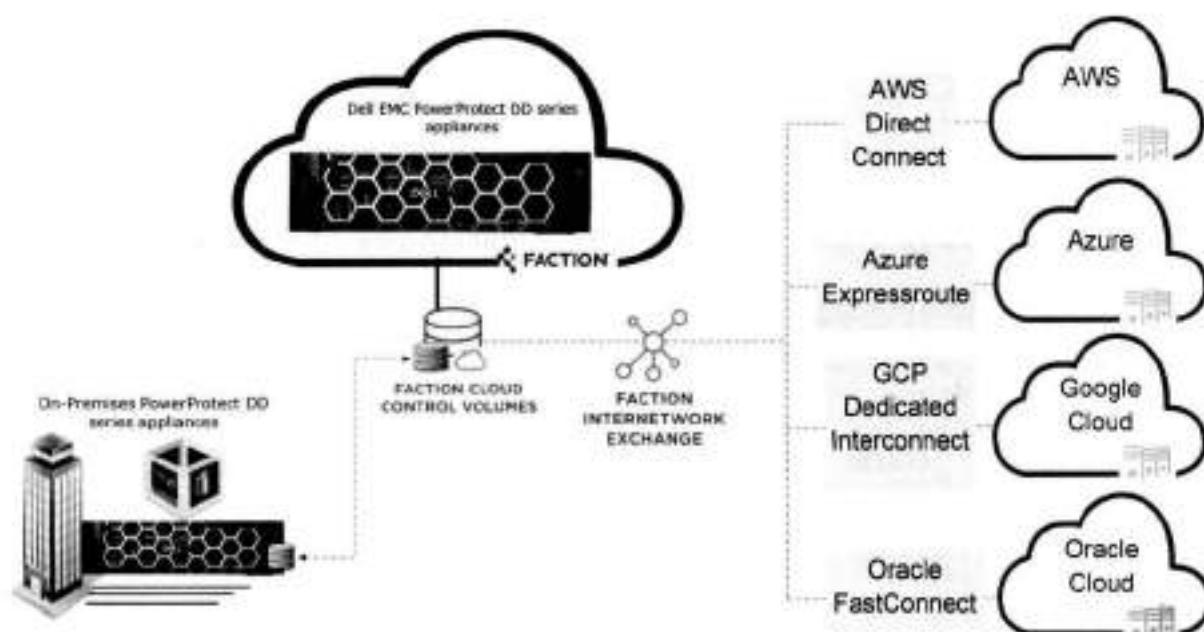
3 Client connectivity

Faction can terminate both Fibre Channel and copper cross-connects in supported facilities, and most other common connections. There are two replication transport options to move data from a customer on-premises data center to the Faction data center:

VPN: Faction can supply an Internet endpoint for replication and client network connectivity over VPN. Also, Faction can terminate IPsec VPNs from compatible equipment for encryption in transit. The VPN must be managed by Faction if the customer does not have a compute environment in the Faction cloud.

Dedicated circuit: Large-scale customers can opt for a dedicated connection for replication traffic between their facility and Faction. These customers can use a VPN temporarily since lead times for dedicated circuits can be in the 90+ day range. Customers may also use a VPN for redundancy to a dedicated link. Faction can source and manage the dedicated link, or the client can work with their carrier directly.

It is the customer's responsibility to manage the network between their on-premises data center and the Faction data center.



For public clouds, customers can establish private connectivity between the respective cloud providers and Faction data center through the following:

- AWS Direct Connect
- Azure ExpressRoute
- Google Cloud Platform Dedicated Interconnect
- Oracle FastConnect

4 Example scenarios

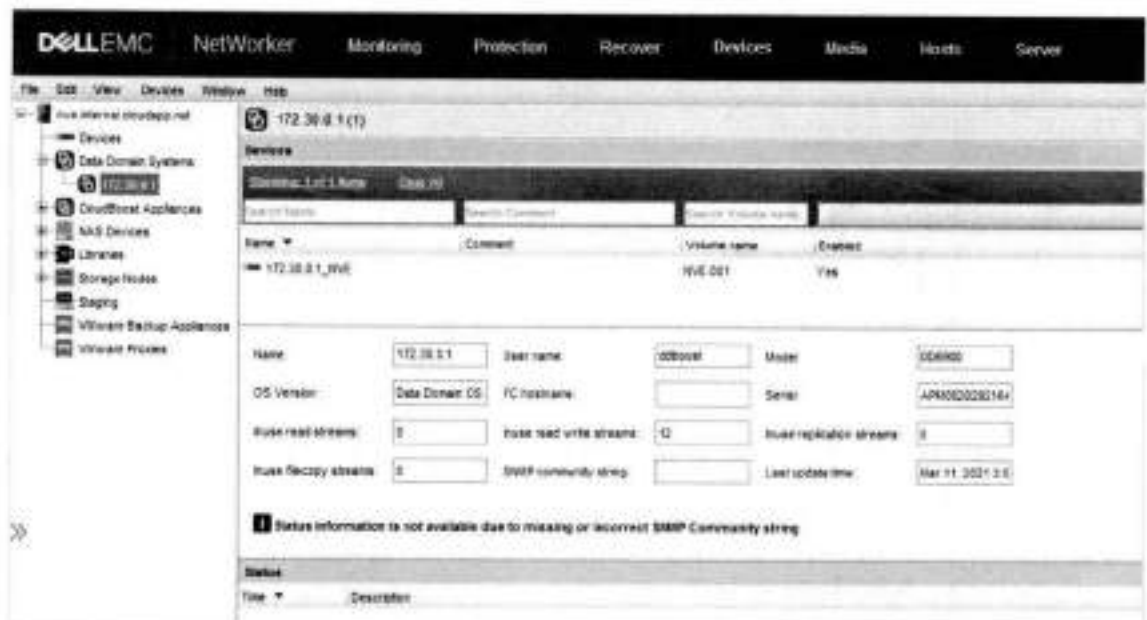
4.1 DD series appliance at Faction as the primary backup target for cloud workloads

In this example, NetWorker Virtual Edition (NVE) is used to back up and restore Azure workloads. A DD series appliance at a Faction data center (DC) is used as the primary backup target.

1. Integrate a DD series appliance at a Faction DC with NetWorker using a ddboost user.



2. Create and configure a new device with the DD series appliance at the Faction DC.



4.2 DD series appliance at Faction as the replication backup target for DDVE deployed on public cloud

In this example, a PowerProtect Data Manager deployed in AWS is used to back up and replicate the AWS workloads. DDVE is deployed in AWS as the primary backup target, and the DD series appliance at the Faction DC is the replication target.

1. Deploy DDVE in AWS, and add a DD series appliance at the Faction DC as Protection Storage in PowerProtect Data Manager.



Example scenarios

2. Create a protection policy, select the DDVE deployed in AWS as the primary backup target, and select the DD series appliance at the Faction DC as the replication target.



3. The policy runs, and the backup and replication jobs complete successfully.



Example scenarios



4. The primary backup copy is available on the DDVE at AWS.



5. The replicated backup copy is available on the DD series appliance at the Faction DC.



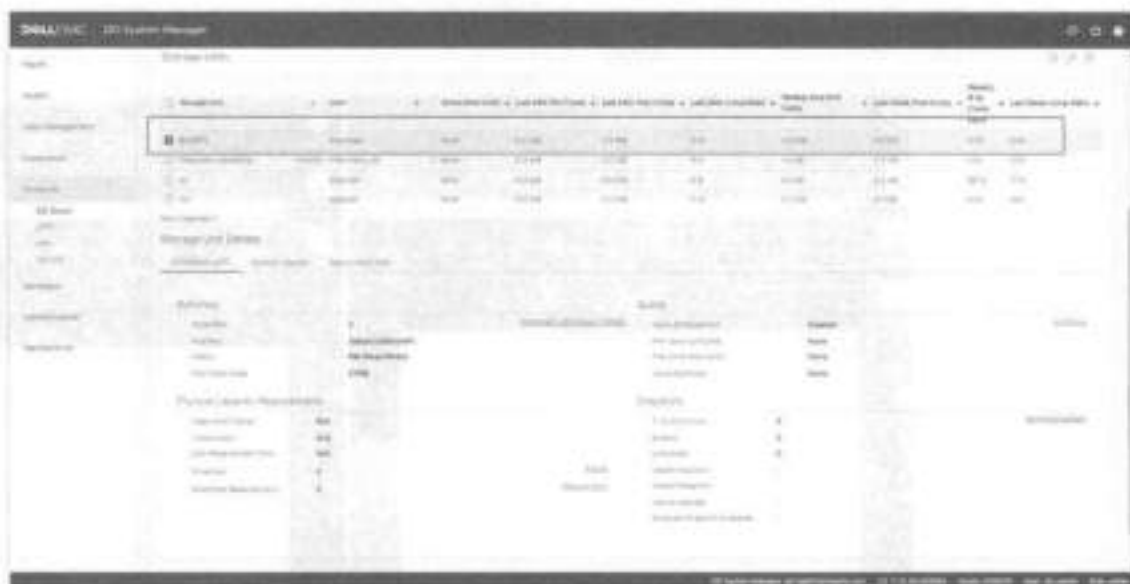
4.3 Mount Faction DD series storage unit on VMs running in multiple cloud environments using BoostFS

In this example, a Storage Unit is created on the DD series appliance at the Faction DC, and it is mounted simultaneously on AWS and Azure virtual machines using the BoostFS plug-in.

1. As a prerequisite, install the BoostFS plug-in, and configure the lockbox on the virtual machines.

Example scenarios


2. Create a Storage Unit on the DD series appliance at the Faction DC.



3. Using the `boostfs mount` command, mount the DD series Storage Unit on the virtual machines.

AWS

Azure

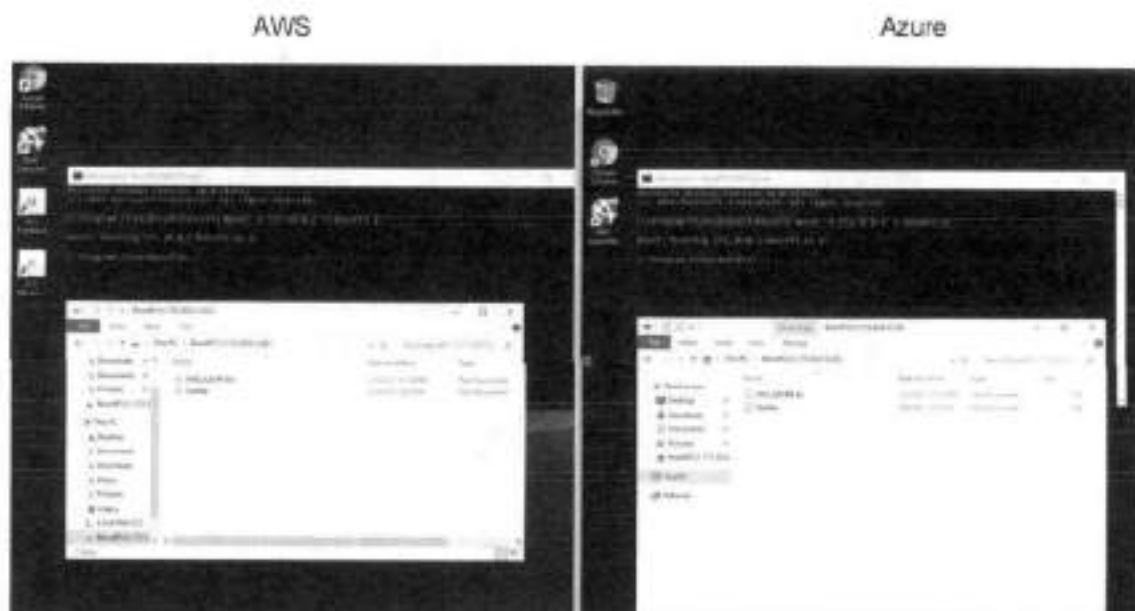


Example scenarios

- The Storage Unit created on the DD series appliance at the Faction DC is successfully mounted on the AWS and Azure VMs.



- The same set of files is available on both VMs using the BoostFS plug-in.



5 Onboarding process and support

5.1 Implementation

The Service Order Form (SOF) or the Statement of Work (SOW) contains the detailed implementation process. Implementation includes the following high-level steps:

1. Solution design and proposal
2. Service order and MSA are signed
3. Introduction and project kickoff meeting
4. Solution design audit and project timeline creation
5. Ready-for-implementation signoff, and install timeline is committed
6. Solution build begins
7. Quality assurance testing and validation
8. Client access provided
9. User acceptance testing (UAT)
10. Close project: Handoff meeting with 24x7 Faction Support team

5.2 Support

The Faction NOC provides first-call support for client circuit issues. If Faction staff cannot resolve the issue, it may be escalated to the vendor for further support and troubleshooting. While Faction remains engaged, clients should be prepared and have resources available to work directly with the vendors to resolve issues.

Description	Client	Faction
Client premise connectivity (Layer 1 beyond the carrier demark)	Primary	
Client premise logical connectivity	Primary	Secondary
Faction premise connectivity		Primary
Faction premise logical connectivity	Primary	Secondary
First-call support		Primary
Middle-mile troubleshooting	Primary	Secondary
Maintenance notification (pass-through from provider only)		Primary

A Technical support and resources

[Dell.com/support](https://www.dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage and data protection technical documents and videos](#) provide expertise to ensure customer success with Dell EMC storage and data protection products.

A.1 Related resources

Access to the following documents may require login credentials.

- [Multi-Cloud Data Services for Dell EMC PowerProtect](#)
- [Data Protection Multi-Cloud Innovations](#)
- [Multi-Cloud Data Services for Dell EMC PowerProtect \(solution brief\)](#)
- [Cyber Recovery with Multi-Cloud Data Services for Dell EMC PowerProtect](#)
- [Multi-cloud Data Protection](#)
- [Data Protection Solution powered by Cloud Control Volumes \(CCV\)](#)
- [Faction CCV](#)
- [Dell EMC PowerProtect DD Series Appliances](#)
- [Dell EMC PowerProtect DD Series Appliances \(spec sheet\)](#)
- [Dell EMC PowerProtect DD Series Appliances \(data sheet\)](#)
- [Dell EMC Data Protection Suite](#)
- [Dell EMC Cloud Disaster Recovery](#)
- [Dell EMC PowerProtect Cyber Recovery Solution](#)

Dell PowerProtect Data Manager: Dynamic NAS Protection

July 2023

H18883.4

White Paper

Abstract

This white paper describes how the PowerProtect Data Manager NAS protection solution protects NAS storage arrays and generic NAS shares.

Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2021–2023 Dell Inc. or its subsidiaries. Published in the USA July 2023 H18883.4.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Contents

Executive summary	4
Introduction	6
Architecture overview.....	8
Protecting NAS assets.....	13
NAS backup workflow.....	29
NAS restore workflow.....	30
Restoring NAS assets.....	31
Performance results	34
References.....	36

Executive summary

Overview

Network attached storage (NAS) is an IP-based file-sharing storage device that is attached to a local area network (LAN). NAS can serve various clients and servers over an IP network. A NAS device uses its own operating system and integrated hardware and software to deliver a range of file-service needs.

NAS is widely used for its simplicity, ease of use, and outstanding performance. With its simple use come challenges regarding data protection. For years, NAS and backup vendors have used the NDMP protocol to protect NAS data. The NDMP protocol has its own limitations, such as manual slicing of a NAS share to achieve multistream backup, limited number of parallel streams, and periodic full backups. Customers also face challenges to protect their growing amounts of data and to back up this data within their specified backup windows.

PowerProtect Data Manager for NAS protection addresses customer challenges of protecting evolving NAS environments. Unlike NDMP-based solutions, dynamic NAS protection is a NAS-vendor-agnostic solution. With dynamic NAS protection, customers can overcome the challenges with the NDMP protocol.

Protecting NAS assets with Data Manager is a non-NDMP solution. Dynamic NAS protection uses the NAS Protection Engine for backup and recovery orchestration. This solution is easy to use and provides automatic discovery, orchestration, and management through the Data Manager UI. With its snapshot technology and intelligent slicing, Data Manager protects NAS data efficiently within the required backup window.

This solution addresses some of the challenges to dynamic NAS protection with the following capabilities:

- Vendor-agnostic solution for NAS protection
- Forever incremental backup and no periodic full
- High number of parallel streams and multiple virtual containers to address scale and performance
- Index, search, and restore
- Restore to any NAS device, such as NFS or CIFS

Audience

This white paper is intended for Dell Technologies customers, partners, and employees who want to use Data Manager to protect NAS storage arrays.

Revisions

Date	Part number/ revision	Description
September 2021	H18883	Initial release
July 2022	H18883.1	19.11 updates
October 2022	H18883.2	19.12 updates and enhancements
April 2023	H18883.3	19.13 updates and enhancements
July 2023	H18883.4	19.14 updates and enhancements: Configuring parameters for asset-level protection

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

Author: Breno Lourdu

Note: For links to other documentation for this topic, see the [PowerProtect Data Manager Info Hub](#).

Introduction

PowerProtect Data Manager for NAS overview

PowerProtect Data Manager for NAS protection is a software-only solution that supports centralized backup and recovery for NAS assets. Dynamic NAS protection provides a non-NDMP, crawl- and backup-based solution by leveraging the NAS Protection Engine internally using Files System Agents (FSA) file-based backup (FBB) technology. Data Manager for NAS protection supports multistream backup and restore. With centralized support, Data Manager controls and manages end-to-end backup and recovery operations.

Data Manager for NAS protection supports all the Data Manager objectives such as DD Replication, Cloud Tier, progress monitoring, and SLA compliance.

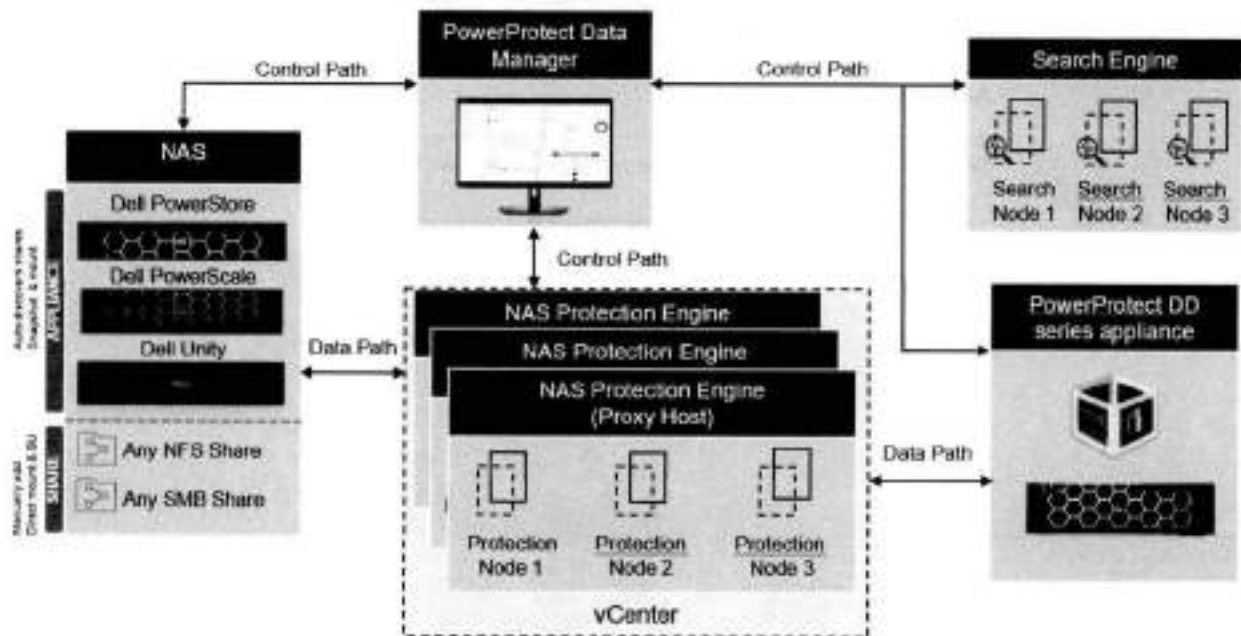


Figure 1. Data Manager for NAS overview

The dynamic NAS solution supports protection for Dell PowerStore, Dell Unity, Dell PowerScale (Isilon) NAS products, and any NFS or CIFS share using generic NAS, such as NetApp, Windows, and Linux file servers, from other vendors.

Note: Supported hardware or software platforms may be updated in subsequent releases. See the support matrix at <https://elabnavigator.dell.com/elh/modernHomeDataProtection> for the latest product information.

Data Manager can protect NAS assets in two ways:

- Appliances: Automatic discovery of shares on supported PowerStore, Dell Unity, and PowerScale (Isilon) products.
- Shares: Network File System (NFS) and Common Internet File System (CIFS) shares from other NAS platforms.

Shares on the supported Dell Technologies appliances (as listed above) are automatically detected, and NFS and CIFS shares from other NAS platforms can be manually added. NAS protection backs up and recovers ACLs and extended attributes for NFSv4 and CIFS or SMB shares.

Data Manager for NAS protection supports the following restore use cases:

- Share-level restore
- Restore to NFS or CIFS shares of the same NAS array or different NAS array
- Restore to original and alternate NAS shares
- File-Level Recovery (FLR): NAS backups are indexed on the Search Engine for search and restore operations

Dynamic NAS protection features

The features of dynamic NAS protection are:

- Software-only solution
 - Data Mover NAS Protection Engine as Virtual Machine with containerized NAS agents
 - Protection Engine OVA
- Auto discovery of NAS shares for supported Dell NAS appliances
- Intelligent share slicer for parallel backups
- Auto-slicing of NAS share to achieve multiple parallel streams
- Auto distribution of backup streams to single or multiple NAS Protection Engine
- Crawl and backup appliance snapshot or generic NAS share
- Forever incremental backup
- Supports large number of parallel backup streams:
 - From 1 to 256 streams (can be set at asset level. See
 - Protection Engine parameters for more details about the scalability of stream counts.)
- Restore with parallel streams:
 - Eight streams (See
 - Protection Engine parameters for more details about the scalability of stream counts.)
- Index, search, and FLR
- Support for multiple protocol shares (discovers two shares: one CIFS and one NFS)
- Beginning with Data Manager 19.11, backups completed with an exception have:
 - A complete list of skipped elements in backup and ACL backup skipped. This information can be generated in the backup log by setting the log debug level to 9. The backup log can then be exported to list the number of skipped files/folders with the reason that they were skipped.

Note: For details about how to review protection logs, see the *PowerProtect Data Manager for Network Attached Storage User Guide* on Dell Support at [PowerProtect Data Manager Info Hub: Product Documents and Information](#).

- Beginning with Data Manager 19.12, backup of Japanese and Chinese file and path names is supported. These languages generally contain file names with multibyte characters and lengths of more than 255 bytes.

Architecture overview

Figure 2 illustrates the dynamic NAS asset backup and recovery solution with Data Manager.

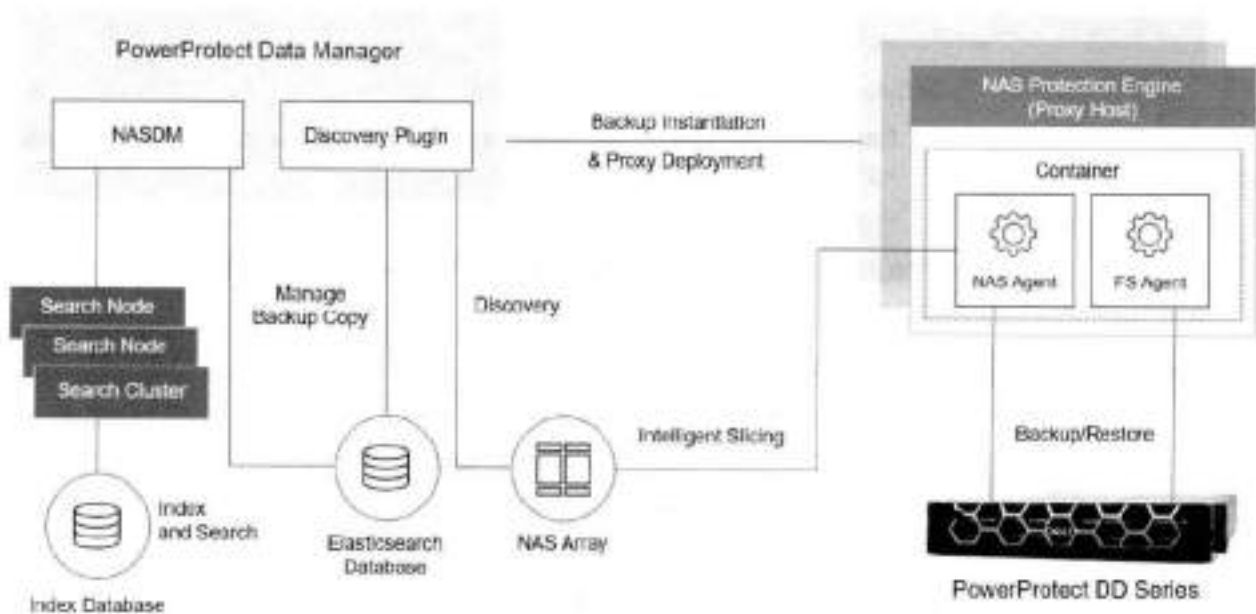


Figure 2. High-level architecture overview

The NAS array is the primary storage location for the NAS data from where the data is read and sent to the secondary storage location, which is PowerProtect DD series appliances. Data Manager uses specialized NAS Protection Engines to protect the NAS assets.

The NAS Protection Engine is used as a data mover for backup and recovery. Containerized NAS agents run on the NAS Protection Engine to support multiple NAS protection operations. Each NAS container is preinstalled with NAS agent and FSA agent. NAS Agent will use FSA binaries for backup and recovery and orchestrate them to run backup and recovery with multiple threads or streams to achieve optimal scale and performance. Once the backup is completed, the Search Cluster creates indexes for the NAS backup and supports search and FLR.

Solution components

Data Manager components for NAS

NAS Data Manager (NASDM): The NASDM microservice in Data Manager is the pillar for NAS workload backup and recovery.

- Orchestrates NAS backup and recovery operations
- Performs NAS backup copy management
- Maintains Protection Copy Set (PCS) and Protection Copy in Elasticsearch database
- Initiates index and search with Search Engine
- Support for DD Replication, Cloud Tier, and Telemetry

NAS discovery plug-in: The discovery plug-in enables automated discovery of supported Dell Technologies NAS appliances as assets and stores the asset information to the Elasticsearch database.

Virtual Proxy Orchestrator Daemon (vpod): NASDM uses vpod to orchestrate NAS assets protection using NAS Protection Engines. NASDM integrates with vpod component for backup and recovery.

NAS Protection Engine components

The Protection Engine is an external virtual machine running on a VMware vCenter for NAS backup and recovery. Each Protection Engine can run multiple containers and each container is preinstalled with NAS agent and FSA agent. A Protection Engine can receive multiple NAS backup and recovery jobs for different shares, and a separate container runs for each job.

NAS Container: Docker container running on Protection Engine for NAS data protection.

- Docker-based container for NAS workloads on the Protection Engine
- Runs on-demand on Protection Engine and destroyed after backup/recovery job completion
- Maximum of three jobs per NAS container will run in parallel
- Multiple containers can run simultaneously
- NAS container is packaged with NAS Protection Engine

NAS Agent: Backup and recovery agent for NAS shares.

- Manages NAS asset snapshot (create, delete, and mount)
- Uses intelligent share slicer to create slices of NAS asset for parallel backups
- Performs multiple stream backup to achieve optimal scale and performance
- Manages Filesystem Agent for backup and recovery
- Manages NAS metadata records for each NAS asset
- Periodically collects backup and recovery progress

FSA Agent: Dynamic NAS Protection uses FSA-FBB as data mover for NAS data protection. FBB method uses a crawl-and-backup method, which means that it crawls the given file system and backs up files and metadata.

- Packaged with NAS agent and installed as part of NAS container deployment on Protection Engine
- Moves the NAS data to and from PowerProtect DD series during backup and recovery respectively

Search cluster

The search cluster consists of one or more Search Engine nodes to index the NAS backup data.

- PowerProtect Data Manager search cluster is a multinode (up to 5) search engine

Note: Each search node can index up to 1 billion files. If there are no virtual machine backups configured and only NAS shares are being protected, each node can index up to 1 billion files over single (or different) shares.

- Deployed as a virtual machine
- NAS protection solution provides file index, search, and FLR
- Allows user to search file or folder across the backups by using the File Search option
- Search is based on file name, folder name, size, time, protocol.

Intelligent auto slicer for NAS protection

The NAS file-share auto slicer is a new library that is embedded in the Data Manager NAS agent. The slicer splits NAS assets (NAS share, a file system) into multiple subassets in preparation for multistream data movement to PowerProtect DD series. Slices are created using parallel threads, and each slice is backed up concurrently using available NAS Protection Engine containers and moved to a PowerProtect DD series appliance.

The slicer partitions NAS assets dynamically before each backup. Based on backup history and changes in the content of the NAS asset being sliced, relevant slices are added, removed, or rebalanced. Periodically, unbalanced trees are automatically managed as content changes over time. No manual reconfiguration is required. The default slice size is 200 GB or 1 million files (tolerance of 30%).

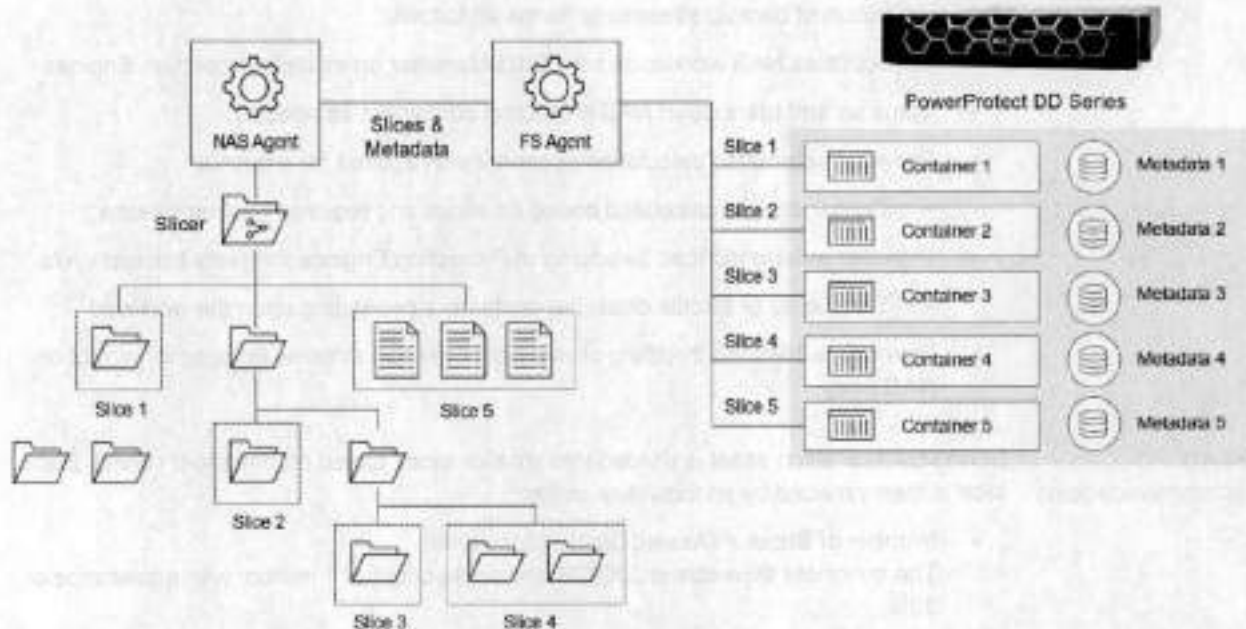


Figure 3. Overview of intelligent auto slicing

Full and incremental behavior:

- Full backup slices: A complete share is traversed in parallel to create slices.
- Incremental slices: Only modified slices are traversed to add or delete slices.
- Dynamic reslicing, rebalancing, and consolidation of slices is based on backup history.

Auto distribution of backup streams

Dynamic NAS solution enables automated load balancing of protection engine hosts, and automatic scaling for containers to achieve maximum backup streams and reduce manual management overhead.




Figure 4. Overview of autodistribution of backup streams

Auto-distribution of backup streams performs as follows:

- Orchestrates NAS workloads from Data Manager on multiple Protection Engines
- Spins up and tears down NAS protection containers as needed
- Provides automated calculation of containers required for a backup
 - Containers are calculated based on slices and required backup window
- Provides automated load balancing of Protection Engines for every backup cycle
 - Throttle up or throttle down the containers depending upon the workload
- Provides automated throttling or number of parallel streams for optimal impact on NAS array

Sizing recommendations

During backup, each asset is divided into smaller slices based on threshold values. Each slice is then serviced by an individual stream.

- $\text{Number of Slices} = (\text{Assets Size}) / (\text{Slice Size})$
The threshold slice size is 200 GB and/or file count of 1 million, with a tolerance of 30%.
- When determining the number of Protection Engines, using a factor of 1.2-1.5x size of the preceding slice count is recommended.
- Each Protection Engine supports up to 24 concurrent streams.

To achieve optimized throughput, follow these guidelines:

- $\text{Number of Protection Engines} = (\text{Number of Slices}) / 24$ (24 is the total count of streams per Protection Engine, where eight streams a piece are served by a different container).
- With the current Data Manager v19.9 release, the recommendation is to scale up to 11 Protection Engines for larger shares (for example 50 TB or larger).
- To achieve optimum performance, it is recommended to use a dedicated 10 GbE network per Protection Engine. The Protection Engine throughput is bounded by underlying network stack on ESXi host. Hence, a dedicated 10 GbE network per Protection Engine would achieve better performance.
- Multiple NAS Protection Engines with a dedicated 10 GbE network can achieve better net aggregated throughput. This includes reading the data from NAS array and writing it to protection storage.
- If the whole environment is 10 GbE network (NAS array, PowerProtect DD series and multiple Protection Engines), the overall throughput is bound by 10 GbE network speed.
- If the read throughput from the NAS array and write throughput to PowerProtect DD series causes a bottleneck with multiple Protection Engines, it is recommended to have more network ports on NAS array and on PowerProtect DD series.
- The asset parallelism per asset helps to load balance the number of streams across multiple shares. Asset Level Parallelism enables all asset backups to run in parallel, and each of these assets has many concurrent streams (as per user input of asset parallelism). Also, if there are enough containers available, all these assets will run

in parallel. The Asset Level Parallelism parameter maximum-supported count with Data Manager v19.9 release is 256 concurrent streams per asset.

- Also, we can use the sizing tool created by Dell Technologies to determine the number of protection engines which must be deployed for a certain protection load. Contact Dell Technologies support to download and use a recent version of the tool. The yellow highlighted sections in the tool can be edited to give us the approx. number of proxy engines which need to be deployed without any manual calculation. Some of the inputs required from the customer are:
 - Array type (PowerScale, Dell Unity, PowerStore, or generic).
 - Total amount of NAS data to be protected in TB.
 - Total number of files in millions.
 - Expected backup duration for full (Gen 0) backup.
 - Network parameters such as the number of array nodes/ports, PowerProtect DD series ports, and proxy ESXi ports involved in the backup. The same data will be used to account for any bottlenecks in the proxy engine calculation.
 - Expected backup duration for synthetic full/incremental (Gen 1) and the approximate change rate expected between backups.

Protecting NAS assets

Steps for NAS protection

The following objectives are required to be completed for protecting the NAS assets:

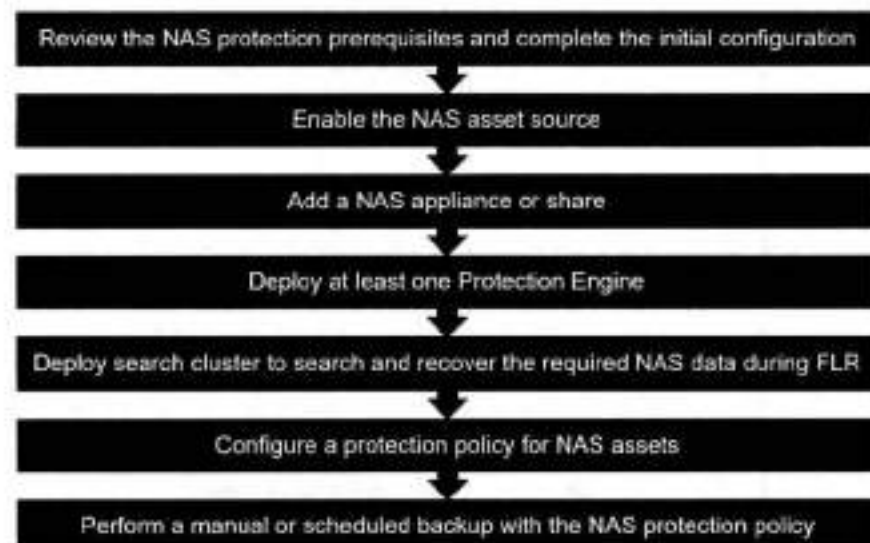


Figure 5. Steps for NAS protection

Note: For details about prerequisites and initial configuration settings, see *PowerProtect Data Manager for Network Attached Storage User Guide* or *release notes* at [PowerProtect Data Manager Info Hub: Product Documents and Information](#) on Dell Support.

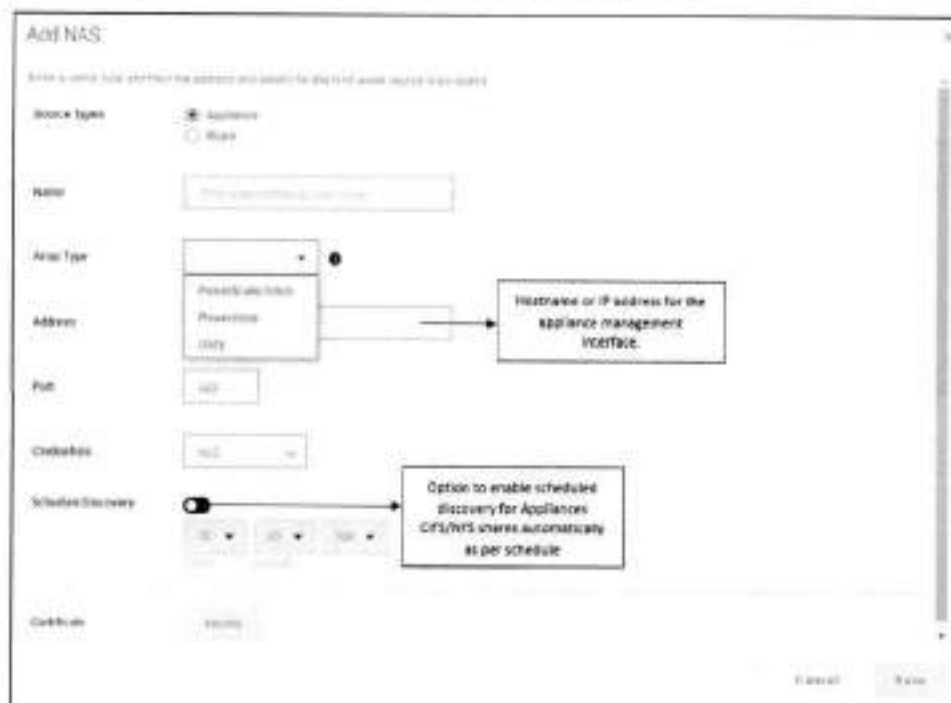
Enabling the NAS asset source

Enabling the asset source in Data Manager allows you to add and register the asset source for the protection of NAS assets. From the Data Manager UI, the NAS asset source can be enabled from New Asset Source section as shown here.



Adding a NAS appliance to Data Manager

For supported appliance types, the NAS appliance can be added as an asset source for Data Manager to automatically discover any assets to protect.



Field	Description
Name	Enter a descriptive name for the appliance.
Array Type	Select a supported appliance type from the drop-down menu (PowerStore, Unity, PowerScale)
Address	Enter the FQDN or IP address for the appliance management interface.
Port	Enter the port number for HTTPS REST API access to the appliance.
Credentials	Select an existing set of management credentials of the NAS array. Alternatively, the Add Credentials option can be selected to provide new credentials. Click Save . The credentials should be root or admin credentials for the NAS array.

Note: Ensure that the production interface is selected for SMB share backups with PowerStore and Dell Unity appliances, whereas NFS backups can use the backup interface. For more information, see the [Dell Powerstore: File Capabilities](#) white paper.

Dell PowerScale SmartConnect and multiple access zone support

With Data Manager 19.12 and later, the following features support Dell PowerScale SmartConnect and multiple access zones:

- PowerScale SmartConnect names can be added as an asset source without duplication of assets during discovery. Thus, we can leverage SmartConnect for client connection load balancing, and dynamic NFS failover and failback of client connections across storage nodes to provide optimal utilization of the cluster resources.
- SmartConnect access zones are used as a data path where backup and restores run across the network, mapped to zones of which the asset is part.
- Shares/exports can be discovered and protected in all PowerScale access zones. Nonsystem access zones can be added as an asset source.

Note: For more information about SmartConnect and nonsystem access zones, see the [PowerScale OneFS Web Administration Guide](#).

Adding a NAS share to Data Manager

For appliances where Data Manager does not support automatic discovery, the NAS share can be added as an asset source.

The screenshot shows a dialog box titled "Add NAS" with the following fields and options:

- Name Type:** Radio buttons for "System" and "Share" (selected).
- Credentials:** A dropdown menu showing "root@nas" with a plus icon.
- Protocol:** Radio buttons for "NFS" and "SMB" (selected).
- NAS Share:** A text input field containing "NAS share URL path" with a plus icon to its right.
- At the bottom right, there are "Cancel" and "Save" buttons.

Table 1. Syntax and port numbers by protocol

Protocol	Syntax	Default Port Numbers
NFS	<NAS>:/<share-path-and-name> <NAS>:<port>/<share-path-and-name>	2049
CIFS	\\<NAS>\<share-path-and-name>	139, 445

Note: <NAS> can be either the fully qualified domain name or IP address for the NAS. Verify and use the user-defined port numbers, if any.

Note: For more detailed information about adding the NAS appliance or share to Data Manager, see the *Dell PowerProtect Data Manager for Network Attached Storage User Guide* at [PowerProtect Data Manager Info Hub: Product Documents and Information](#) on Dell Support.

Discovering the NAS asset sources and assets

Discovered NAS asset sources

When the NAS appliance or share is successfully added to Data Manager, all NAS asset sources are listed on the asset sources section as shown in this section. By default, initial discovery is done automatically when the NAS asset source is added and when subsequent discoveries are either manual or scheduled. Using the Schedule Discovery option (see the section *Adding a NAS share to Data Manager*), a full discovery at a certain time every day can be scheduled on a given time. At any time, an on-demand discovery of the assets can be done using the Discover option.



Note: The discovery status of generic NAS asset sources is displayed as Unknown.

Discovered NAS assets

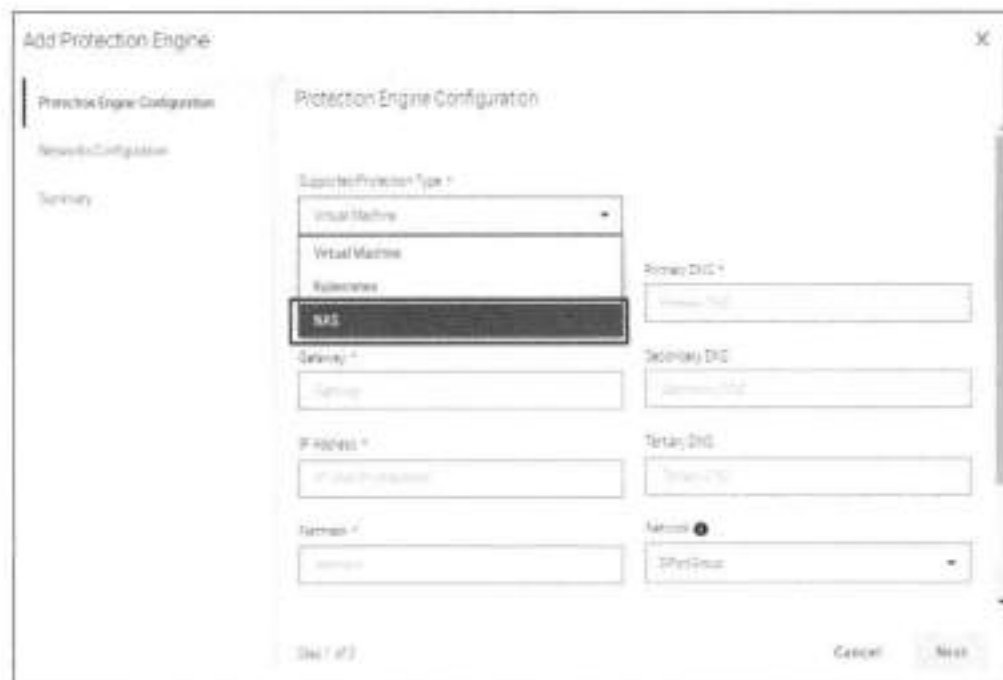
After a successful discovery, a list of the NAS assets that are discovered from the appliance is displayed. The discovery time is based on the network bandwidth and number of assets on the NAS appliance.

**Notes:**

- For generic NAS assets and Dell PowerScale, the size of the share is shown as 0 bytes during the initial discovery. However, the size of the share is determined and displayed after the first successful backup.
- If there are multiprotocol shares in supported Dell appliances, Data Manager displays two entries of the same asset with different protocols (CIFS or NFS).

Deploying a Protection Engine for NAS asset protection

The NAS Protection Engine is deployed on the selected VMware vCenter, and Data Manager registers the Protection Engine. The NAS Protection Engine hosts the NAS agent and FSA.



Protection Engine parameters

A NAS Protection Engine can protect multiple NAS assets simultaneously by hosting separate containerized agents for backup and restore operations on individual shares. Data Manager selects a Protection Engine for each protection job based on availability.

- Each Protection Engine has 24 parallel streams that can be used to protect NAS assets. However, a user-defined count of parallel streams can be specified for a particular NAS asset.
- By default, Data Manager allocates eight streams per NAS asset and can allocate from one to 256 streams per asset, which is tunable. For restore operations, eight streams are allocated, and it is not tunable.
- To set the stream count for a particular NAS asset, go to **Infrastructure > Assets** on the Data Manager UI. From the NAS table, select an asset from the list and then select **More Actions > Set Stream Count**.



Enter a new value (1 to 256) for **Maximum Streams** and click **Save**. The default value is 8 streams.



Centralized protection policy for NAS assets

Data Manager supports centralized protection for NAS assets, where all the stages of the protection policy are managed by Data Manager.



Note: Data Manager uses these credentials at the policy level for all shares unless otherwise specified at the asset level. The credentials provide snapshot creation and export permissions on the appliance and read/write access to the NAS shares.

If credentials are set at the protection policy level, all shares should use the same credentials for access. Otherwise, if individual asset credentials are set at the asset level, multiple assets use their respective credentials.

For more information about creating a protection policy for NAS protection, see the *Dell PowerProtect Data Manager for Network Attached Storage User Guide* at [PowerProtect Data Manager Info Hub: Product Documents and Information](#) on Dell Support.

Individual asset credentials settings

Individual asset credentials can be set from the Data Manager UI.



Configure asset parameters

Beginning with Data Manager 19.14, you can configure various asset-level parameters from the Data Manager UI.

To configure the parameters:

1. Select the asset and then select **More Actions > Assign Configurations**.



2. In the **Configure Asset Parameters** window, turn on the **enable asset level parameters** toggle switch to enable asset-level configuration.



You can then configure these backup and job options:

- Choose to continue with the backup if file or folder names exceed 255 bytes and skip the files and folders whose names exceed that character limit.

The following table describes the behavior when both flags are selected, only one flag is selected, and when no flag is selected:

ContinueOnFileNameLenLimitReached	SkipFilesWithNameLenLimitReached	Expected behavior
TRUE (default)	TRUE (default)	Backup continues and skips the files/folders with names >255 bytes.
TRUE	FALSE	Backup continues and includes the files/folders with names >255 bytes.
FALSE	FALSE	Backup fails when the files/folders with names >255 bytes are encountered in the share.
FALSE	TRUE	Backup fails when the files/folders with names >255 bytes are encountered in the share.

- Specify the number of times that Data Manager should attempt automatic retry of failed slices. Accepted values are from 0 to 5.
- Specify a timeout interval after which Data Manager will cancel an unresponsive backup/recover session. Values can range from 3 hours to 12 hours.

Enable indexing and selecting backup behavior for file exceptions

When the Search Engine is installed, the Enable indexing for file search and restore option is available while creating the protection policy to enable indexing the NAS backups.

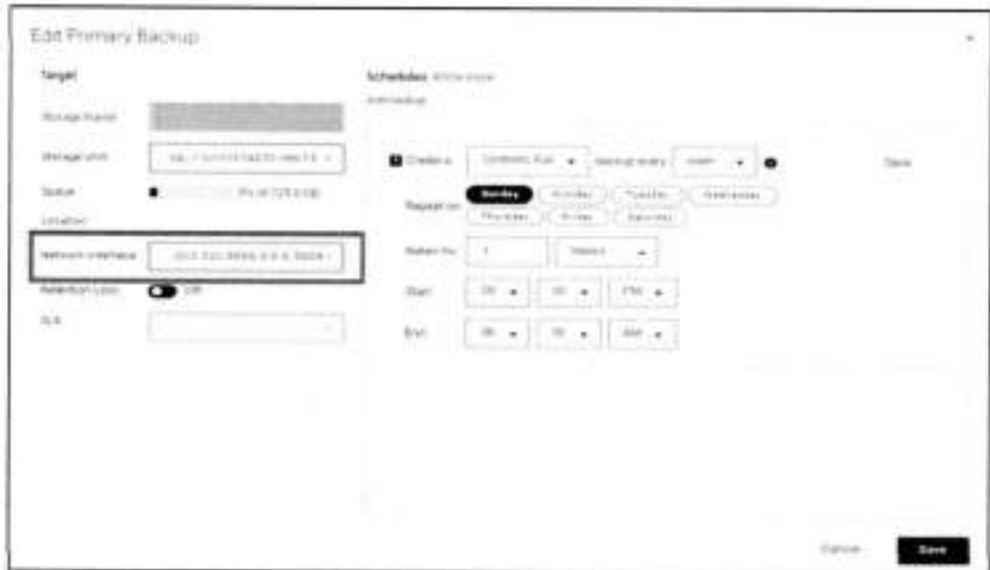
Beginning with Data Manager 19.12, you can continue the backup even if a data access denied or an ACL access denied failure is encountered on files. You can set these flags as shown here:



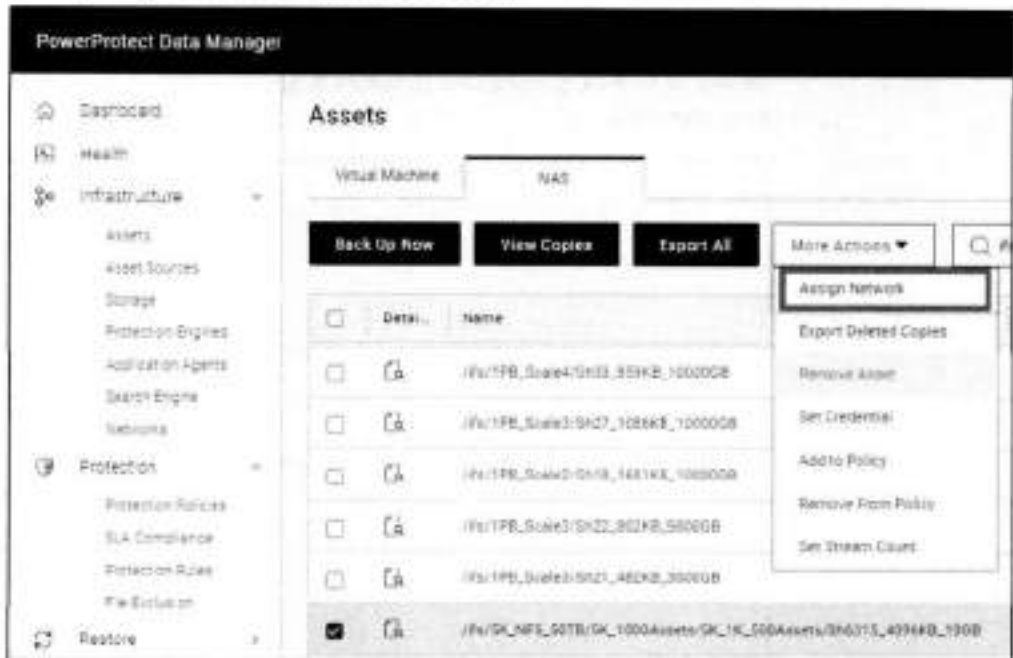
IPv6 support

Beginning with Data Manager 19.13, IPv6 networks can be used for backup, restore, and FLR operations. IPv6 can be selected by using one of the following methods:

- Assigning the network interface at the policy level. This method will assign the IPv6 network interface to all the underlying assets.



- Assigning the network interface at the asset level. This method will assign the IPv6 network interface to an individual asset and takes a higher precedence than the policy level network interface selection.





Note: To perform any backup, restore, or FLR operations, the management IP, NAS server IP, and the DD network interface should belong to the same network family.

Multi VLAN support

Beginning with Data Manager 19.13, NAS workloads support multi-VLAN configuration for management and data transfer operations such as backup, restore, and FLR. Having separate VLANs configured in the environment ensures that the data transfer happens through dedicated VLANs.

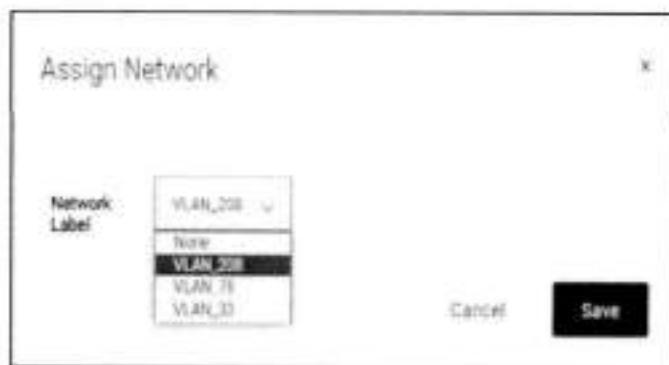
Use the following methods to select the VLAN during the proxy deployment:

- Select the VLAN ID at the policy level. This method assigns the network interface to all underlying assets.



- Select the VLAN ID at the asset level. This method assigns the network interface to an individual asset and takes a higher precedence than the policy level network interface selection.

Protecting NAS assets



Note: During proxy deployment, the default VLAN ID 0 will be tagged to NAS proxies.

Manual protection of NAS backup

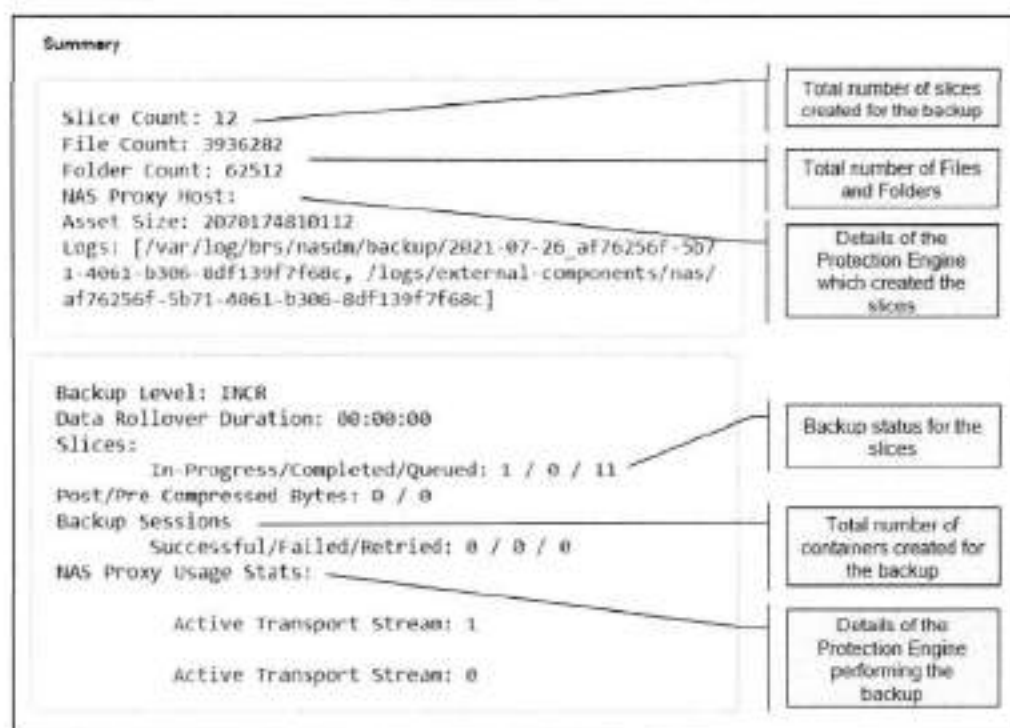
With the Protect Now option, you can select Synthetic Full or Full as the backup type.



Jobs > Protection Jobs shows the protection job details in the Data Manager UI for the NAS protection job in progress:



The following figure shows the job summary details in the Data Manager UI for the NAS backup in progress:



Jobs > Protection Jobs shows the protection job details in the Data Manager UI for the successful NAS backup:

Protecting NAS assets



The following figure shows the job summary details in the Data Manager UI for the successful NAS backup:



Backup completed with exceptions

Beginning with Data Manager 19.12, the skipped element count of ACLs and backed-up data is captured in the job summary:



The user can see all skipped elements in the UI, using the **Show Skipped files** button, for index enabled backups.



Existing Export Log functionality is used to download the skipped element list in CSV format.

File Exclusion filter

Beginning with version Data Manager version 19.13, you can configure filters to exclude files and folders from the backup. You can specify one of the following filter options:

- Type of file
- Size of file

Protecting NAS assets

- Modified time
- Folder path

We can either create the filters during policy creation or create it from Protection > File Exclusions in the Data Manager UI as shown here:



Excluded File/Folder Count is captured in the job summary.

```
Backup level: full
Data Rollover Duration: 01:05:18
Slices:
  In-Progress/Completed/Queued: 0 / 406 / 0
Post/Pre Compressed Bytes: 1264296 / 3479416
Protected Bytes: 1002
File/Folder Count: 2 / 3906
Transaction id:1679664670290447904
Backup Sessions:
  Successful/Failed/Retried: 51 / 0 / 0
Excluded File/Folder Count: 78123949 / 0
```

NAS backup workflow

The NAS agent backup workflow is a three-step procedure, and it supports only a centralized backup workflow. Data Manager initiates the NAS-share backup through a protection policy. The NASDM microservice running on Data Manager initiates the NAS-share backup request using the Virtual Proxy Orchestrator Daemon (vpod). Backup requests are received on the Protection Engine and are routed to the NAS agent. Depending upon the array support, the NAS agent creates the snapshot and uses the slicer to create slices of NAS data. The file system agent moves the data slices in parallel to the PowerProtect DD series appliance. NASDM initiates the indexing when the backup is complete.

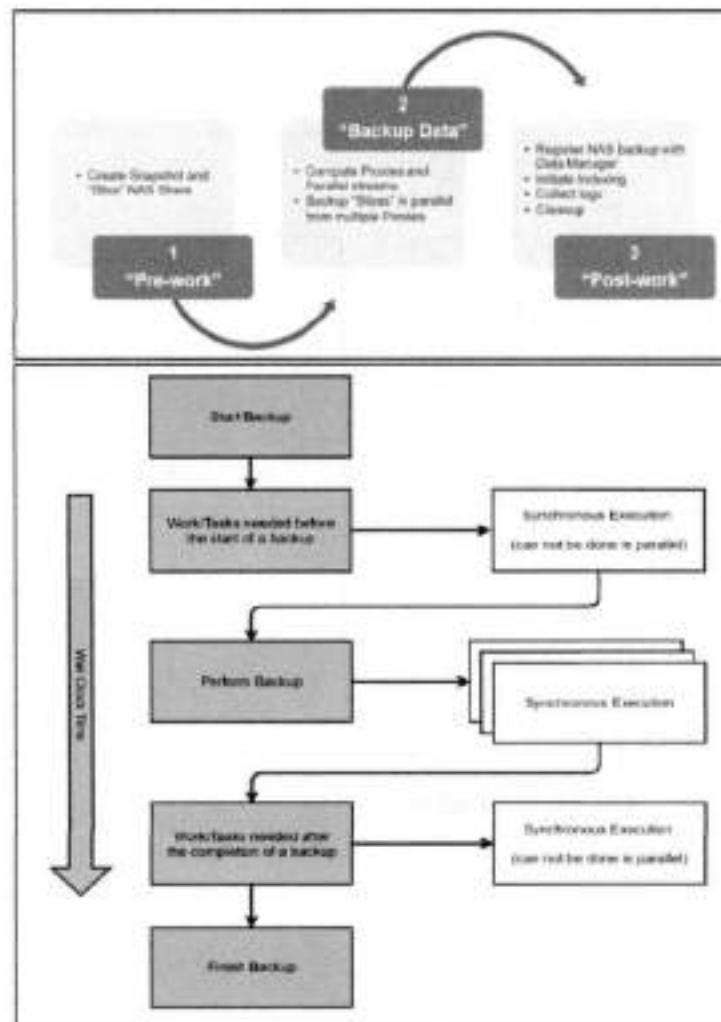


Figure 6. NAS backup workflow

NAS restore workflow

The NAS agent recovery is a single-step procedure, and it supports only a centralized restore workflow. The recovery is orchestrated using NASDM and Protection Engine infrastructure. Upon user selection for recovery (share level or search and restore), the recovery is initiated from Data Manager.

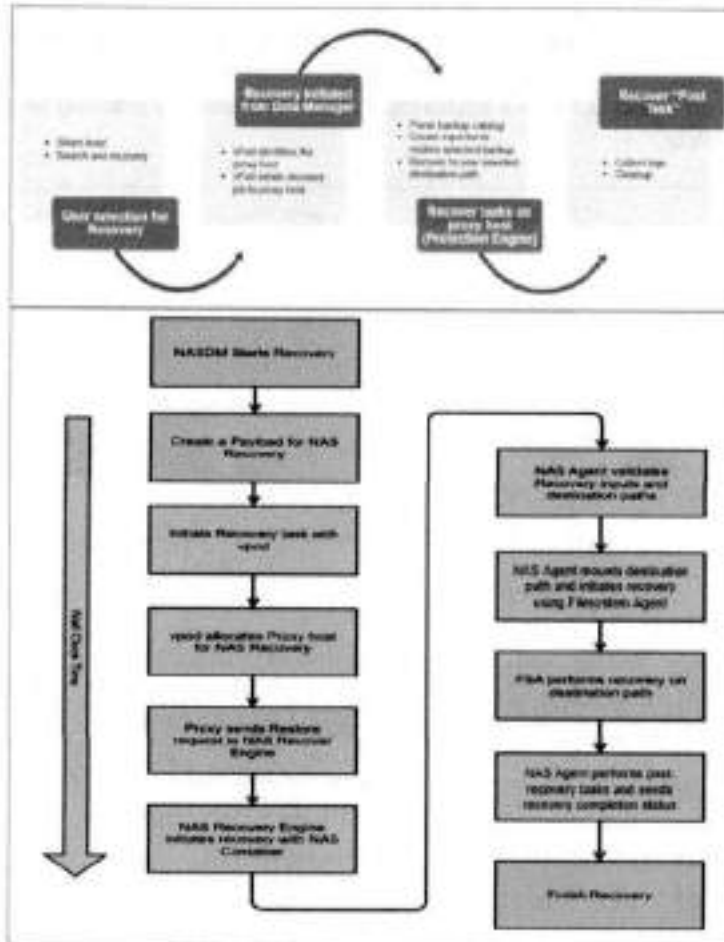
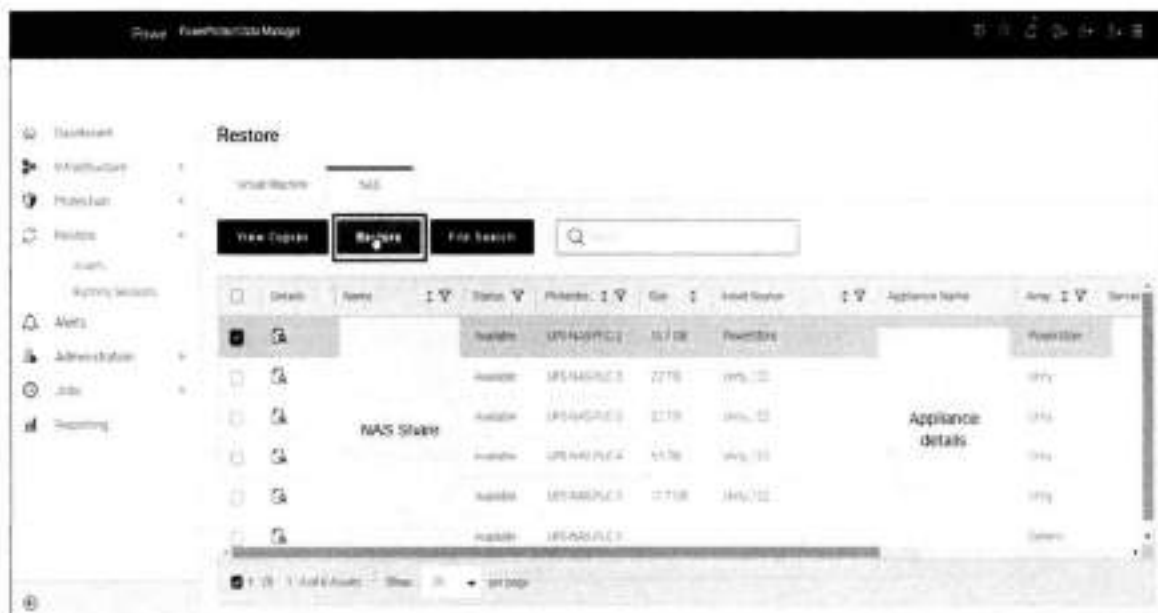


Figure 7. NAS restore workflow

Restoring NAS assets

Data Manager provides support to restore a NAS asset to the original location or to an alternate location. Data Manager also supports FLR using the search engine to restore individual files and folders from NAS backups.



The **Location** pane shows the options to restore and overwrite the original share or restore to an alternate share or array:



If you are restoring to an alternate location, a table of available shares is displayed. Review any warnings and click **OK**. The yellow warning on the shares indicates the presence of a warning for the destination share when the destination free space is less than the size of the restore data.

NAS file-level restore using File Search

The NAS protection solution provides File Index, Search, and FLR, which allow you to search the file and folder from the entire NAS backup. When the Search Engine is deployed and the NAS protection policy is enabled with indexing, you can use the File Search option to restore individual files and folders from one or more NAS backups.

Restoring NAS assets



Beginning with Data Manager 19.12, you can search for skipped files that have not been backed up by using the File Search option.

Note: Indexing must be enabled for the File Search option to become available.



File versions: Select how Data Manager should distinguish files from different backups if the selected files and folders exist in multiple backups of the same NAS asset.



Option	Description
In the Same Folder	The restore appends a suffix to the filename. The suffix identifies the backup from which the file or folder was restored.
In a Separate Folder	The restore uses separate folders to group files from different backups. The folder name identifies the backup from which the file or folder was restored.

NAS Restore > Location shows the destination restore location options for FLR:



Performance results

Option	Description
Restore and Overwrite the Original Files and Folders	The restore operation overwrites any files at the original location with the same names.
Restore to an Alternate Share or Array	A table of available shares appears. Complete the substeps.

Beginning with PowerProtect Data Manager 19.11, having FLR restore to the original location allows you to overwrite the existing files on the export/share without creating the entire path hierarchy of the restored file/folder.

For alternate share restores, you can recover the files/folders directly at the root level of the alternate asset/share/export. This functionality helps to retain the complete folder hierarchy when FLR actions are performed.

Performance results




Figure 8. Backup performance results

Disclaimer: Backups were performed using PowerScale storage as the source and PowerProtect DD series as the backup target with PowerProtect Data Manager. These results were derived from Dell Technologies internal testing performed under varying conditions.

Filesystems come in all sorts of complex topologies. In this section, we observe the average write throughput achieved during backups for file systems with Balanced directory structure and Unbalanced directory structure. Data Manager uses auto slicing and intelligent scaling through proxies to support faster multistream backups irrespective of the file system type.

With a Balanced directory structure (millions of 512 KB files distributed in multiple directories), the initial full backup to PowerProtect DD series is completed with a write

throughput of 4 TB/hr. A subsequent incremental backup (with 4 percent change in data) shows an increased throughput of 5 TB/hr.

Comparable results are seen with an Unbalanced directory structure as well (millions of files varying from 32 KB to 4 GB file sizes distributed in multiple directories). During an initial full backup to PowerProtect DD series, write throughput is 2 TB/hr and the incremental backup (with 4 percent change in data) shows an increased throughput of 3 TB/hr.

For these tests, multiple proxy engines were used to support the workload, and the asset level parallelism was set to 256. To learn more about calculating number of proxies required and setting protection engine parameters, see Sizing recommendations.

References

Dell Technologies documentation

For more information about Dell Technologies data protection solutions, see the [Data Protection Info Hub](#).

The following Dell Support documentation provides information related to this document. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

- PowerProtect Data Manager guides at [PowerProtect Data Manager Info Hub; Product Documents and Information](#):
 - Dell PowerProtect Data Manager for Network Attached Storage User Guide
 - Dell PowerProtect Data Manager Administration and User Guide
 - Dell PowerProtect Data Manager Deployment Guide
 - Dell PowerProtect Data Manager Release Notes
- PowerProtect DDOS documentation for PowerProtect DD series appliances:
 - [Dell DDOS Administration Guide](#)

PowerProtect Data Manager: VMware Virtual Machine Protection Using Transparent Snapshots

June 2023

H18884.6

White Paper

Abstract

This white paper provides insights into how to protect and restore VMware virtual machines using transparent snapshots available with Dell PowerProtect Data Manager.

Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2023 Dell Inc. or its subsidiaries. All Rights Reserved. Published in the USA June 2023 H18884.6.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Contents

Executive summary	4
Introduction	5
Transparent snapshots architecture	5
Integration with PowerProtect Data Manager.....	7
Transparent snapshots life cycle	12
Restoring virtual machines	16
Performance test results	18
References.....	21

Executive summary

Overview

This white paper describes how to protect and restore VMware virtual machines using transparent snapshots available with Dell PowerProtect Data Manager. It details the architecture and life cycle of transparent snapshots and describes how this capability is integrated with PowerProtect Data Manager. It also includes an overview of the process for restoring virtual machines and presents test results that show the performance benefits of this solution.

Revisions

Date	Part number/ revision	Description
September 2021	H18884.1	Initial release
December 2021	H18884.2	Content updates
March 2022	H18884.3	Update for PowerProtect Data Manager 19.10
October 2022	H18884.4	Update for PowerProtect Data Manager 19.11 and 19.12
March 2023	H18884.5	Minor updates
June 2023	H18884.6	Updated to include PowerProtect Data Manager 19.13 enhancements

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

Author: Idan Kentor

Note: For links to other documentation for this topic, see the [PowerProtect Data Manager Info Hub](#).

Introduction

The VMware virtual machine (VM) backup process transfers or exports data from a VM within a VMware environment to a secondary protection storage system. The Dell PowerProtect appliance can be on a primary or secondary site, or in the cloud. A backup engine or software such as PowerProtect Data Manager manages this process. PowerProtect Data Manager can perform data management and copy management operations on the backup copies and ensure that all data is cataloged properly. This function makes available a consistent VM copy as part of a restore requirement in a disaster scenario.

PowerProtect Data Manager can protect VMware VMs in a reliable and efficient manner using VMware vSphere Storage APIs - Data Protection (VADP) snapshots (see Figure 1). These VADP snapshots are reliable, proven, and certified by VMware, and can be used as part of backup operations.

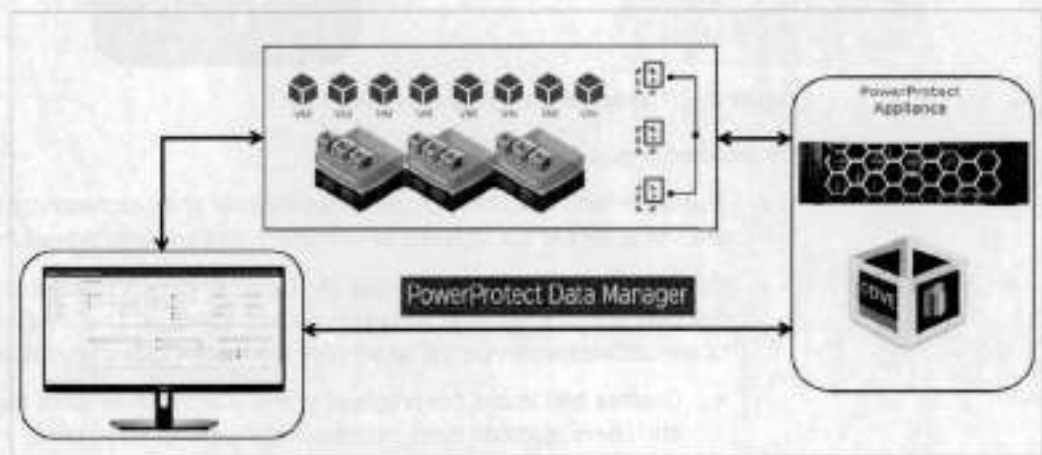


Figure 1. PowerProtect Data Manager for VMware

However, the VADP snapshot process pauses the execution of the VM and allows in-flight disk I/O operations to be completed. This action might increase the read and write latency and affect the snapshot and VM ecosystem life cycle. When the life cycle of a VADP snapshot is analyzed, the snapshot entry and exit points inflict a penalty on a VM. After a snapshot of a VM disk file is produced, requiring the VM to be stunned, a snapshot of the VM disk file is ingested. Then, the deltas must be consolidated into the base disk. When you create a snapshot of a high-transactional application, such as a database, there can be adverse effects. These effects include lengthy backup windows and application timeouts when the stun to ingest and consolidate the workflow is not efficiently managed.

Addressing these issues requires a solution that can deliver not only backup and restore capabilities but also an alternative way to reduce the adverse effects of the VM stun operation.

Transparent snapshots architecture

As shown in Figure 2, PowerProtect Data Manager transparent snapshots use the vSphere API for I/O (VAI/O) Filtering framework. The transparent snapshots data mover (TSDM) is deployed in the VMware ESXi infrastructure through a PowerProtect Data

Manager VIB. This deployment creates consistent VM backup copies and writes the copies to the protection storage (PowerProtect appliance).

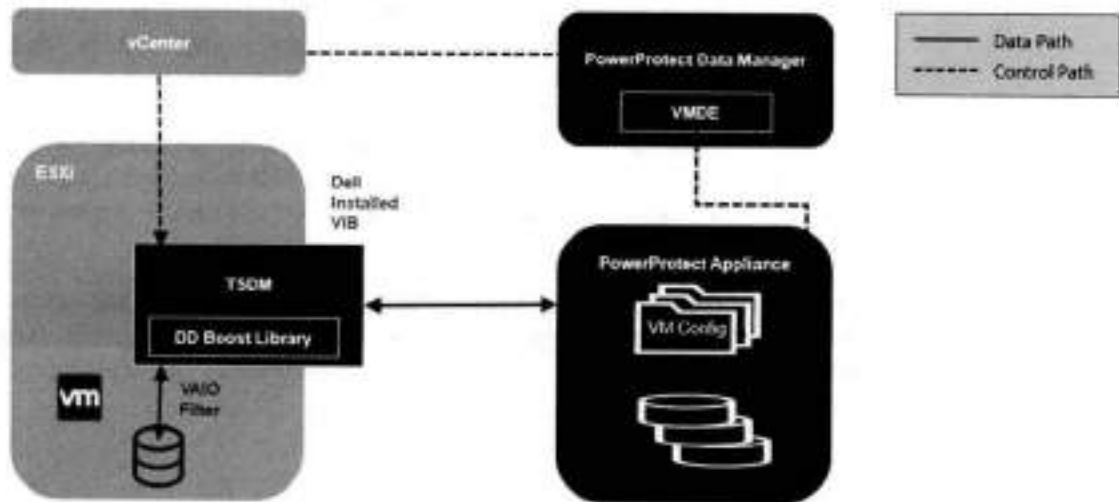


Figure 2. Transparent snapshots architecture

On the control and data paths:

- PowerProtect Data Manager assumes the role of an orchestrator where it identifies the VM assets in the VMware environment and provides scheduling capabilities.
- PowerProtect Data Manager uses VM Direct Engine (VMDE) to communicate with the VMware vCenter level APIs provided by VMware. The VM Direct Engine communicates with vCenter to achieve the following two key tasks:
 - Creates and tracks the progress of the vCenter level tasks that are visible to the end users, such as sync, restore, and snapshot operations
 - Is responsible for locating the relevant ESXi host on which the operation (backup or restore) is to be performed, based on the placement of the VM asset to be protected
- On each ESXi host, the protection-related APIs and workflows from VMware are facilitated using a VAO filter.
- Each ESXi communicates with the Transparent Snapshot Data Mover (TSDM) component, which is responsible for the VM-backup data movement.
- The backup and restore processes transfer the transparent snapshots respectively to and from the PowerProtect appliance.
- TSDM also consists of the PowerProtect appliance SDK (DD Boost library), which helps the framework access the storage units on the PowerProtect appliance. It also helps write and read data from those storage units.

Note: PowerProtect Data Manager manages the TSDM component by using the VIB (VMware Certified) from Dell Technologies. This component is installed dynamically as part of the integration of PowerProtect Data Manager that requires protection of VMs using transparent snapshots. The APIs being used are supported in VMware ESXi 7.0 U3 and later.

Integration with PowerProtect Data Manager

This section examines the steps to integrate PowerProtect Data Manager within the VMware infrastructure, including deploying all necessary components and enabling VM protection using transparent snapshots.




Figure 3. Integration with PowerProtect Data Manager

After VMware VMs are discovered as PowerProtect Data Manager assets, the next steps are creating a protection policy and adding the VM assets for protection.

Note: For information about the criteria for policies to be eligible for transparent snapshots, see [Criteria](#).

1. PowerProtect Data Manager creates the storage unit in the PowerProtect appliance for storing backups.
2. According to the schedule defined in the protection policy, the scheduler activates.
3. As these steps occur, the PowerProtect Data Manager VM Direct Engine initiates API calls to vCenter. It then validates the ESXi version (7.0U3 and later), uploads and installs the PowerProtect Data Manager TSDM VIB, and enables the service. Then, the VAIO filter is attached to each VM disk. In this step, the TSDM component is created but remains idle (running but not used) because there is no data movement. You can see the VIB file installed on the ESXi host that houses the VM being protected (Figure 4).

Library		Files					
VIB Key	: null	<input type="checkbox"/>	DE_postback_tool_593-ADK13550176025.v	429 KB	07/06/23 10:22 PM	File	jeval@DE_06026_jostark
VIB Operation	: (R, install)	<input type="checkbox"/>	netdata.vib	11 KB	07/06/23 10:22 PM	File	jeval@DE_06026_jostark
Vendor	: S						
Vendor Support	: null@net.com						

Figure 4. The VIB file installed

Note: Installation or upgrade of the TSDM VIB does not require the target ESXi host to be in maintenance mode.

The VIB deployment process operates on the relevant ESXi hosts concurrently—25 ESXi hosts at a time.

The process skips ESXi hosts that are powered off or in maintenance mode. The mechanism also has integrated logic to detect and prevent upload of the VIB package to hosts that already have the package in one of their datastores.

4. When vCenter acknowledges success, PowerProtect Data Manager marks the VM to be protected by transparent snapshots.

Note: Aside from the DEL_bootbank VIB file, Figure 4 shows a metadata.zip file that contains information related to the VIB, such as dependencies on the host, system requirements, summary, and version. The files are in the Datastore Files section of the ESXi host. This VIB installation is also shown in the PowerProtect Data Manager Policy Config Job summary.

Criteria

The criteria for policies to be eligible for transparent snapshots are:

- Crash Consistent
- Performance Optimization mode or Capacity Optimization Mode (with PowerProtect Data Manager version 19.10 and later)
- Swap File Exclusion: Disabled
- Quiesce Filesystem: Disabled

Policies created before PowerProtect Data Manager 19.9 will not automatically start to use TSDM upon upgrade to 19.9 and later. The same is true for policies created in 19.9 or earlier with Capacity Optimization mode: TSDM will not be used automatically upon upgrade to PowerProtect Data Manager 19.10. The Data Mover type would be updated the next time the policy is edited or when it is explicitly configured. When a full backup would be performed for the first time, TSDM operates then and when the policy option is switched between Performance and Capacity Modes. TSDM remains the Data Mover when Performance Optimization mode is switched to Capacity Optimization mode, and conversely.

PowerProtect Data Manager 19.10 enhancements

As a precautionary measure, starting with version 19.10, PowerProtect Data Manager uses VADP automatically in the following cases:

- VM with RDM disks
- VM with more than 40 disks
- VM with Fault Tolerance (FT) enabled

Also, with PowerProtect Data Manager 19.10, backups that are taken with TSDM as the data mover can replicate to the cloud using Cloud DR for all supported cloud protection and recovery use cases.

**PowerProtect
Data Manager
19.11
enhancements**

Beginning with version 19.11, the PowerProtect Data Manager UI provides an option to override the automatic protection engine selection and manually select the VM Direct protection engine to be used. It also enables the option to migrate the protection engine being used on an asset basis. For example, VADP is being used to back up a certain VM and now the backup admin wants to leverage TSDM so the asset protection engine can be migrated to TSDM.

Limitations

As a precautionary measure, PowerProtect Data Manager does not support the following with transparent snapshots:

- Physical or virtual RDM disks
- VMs with encrypted VMDKs
- VMs with more than 40 disks
- VMs with Fault Tolerance (FT) enabled
- Azure VMware Solution (AVS) on Microsoft Azure
- Google Cloud VMware Engine (GCVE) on Google Cloud Platform (GCP)
- VMware Cloud (VMC) on Amazon Web Services (AWS)
- VMware Site Recovery Manager (SRM) cannot co-exist with TSDM on the same VMs

Override protection engine

The PowerProtect Data Manager UI provides an option to override the automatic protection engine selection and manually select the VM Direct protection engine to be used. It also enables the option to migrate the protection engine being used on an asset basis. For example, VADP is being used to back up a certain VM and now the backup administrator wants to leverage TSDM so that the asset protection engine can be migrated to TSDM.



Figure 5. Protection mechanism override

PowerProtect Data Manager 19.12 enhancements

The following features and changes were introduced in version 19.12:

1. The number of concurrent TSDM jobs are doubled to 20 for both backup and restore per ESXi host. In previous releases, the maximum backup and restore jobs were 10 for each type. This new throttling mechanism is available with PowerProtect Data Manager 19.12 and vSphere 7.0U3d and later.
2. VIB deployment enhancements:
 - The VIB deployment process now operates on the relevant ESXi hosts concurrently - 25 ESXi hosts at a time.
 - The VIB deployment process skips ESXi hosts that are powered off or in maintenance mode.
 - The VIB deployment process has been enhanced to prevent upload of the VIB package to hosts that already have the package in one of their datastores.

Protection parallelism

With vSphere 7.0U3d and later, there can be a maximum of 20 concurrent TSDM jobs – backups and restores—per ESXi host. The limit is set to 18 concurrent backup jobs or 16 concurrent restore jobs per host. It is a shared pool in which neither backup nor restores can reach 20 so that other restore and backup jobs can run at the same time.

With vSphere 7.0U3c, there is a static limit of 10 concurrent backup jobs and 10 concurrent restore jobs per ESXi host.

There can be maximum of 180 concurrent VM operations per vCenter.

PowerProtect Data Manager 19.13 enhancements

PowerProtect Data Manager 19.13 introduced the following capabilities:

- **Restore storage policies**—The option to assign, upon restore, the VM and its disks to the set of storage policies assigned at the time of the backup.
- **Backup and restore encrypted VMs (VMcrypt)**—Support for protection of encrypted VMs. Support for encrypted VMs depends not only on PowerProtect Data Manager 19.13 but also on vSphere 8.0 (patch b). An encrypted VM is always backed up as unencrypted. If the restore storage policy option is not selected as part of the restore flow, then the VMDKs would be restored as unencrypted, and the default datastore storage policy would be used. The restored VMDKs would be encrypted and assigned to the VM encryption policy if the option to restore storage policy is selected and the original VM disks were assigned to encryption storage policy. In such cases, the encryption would be a post-restore action on the vSphere side once the encryption storage policy gets associated with the restored VMDKs.
- **Restore VM BIOS UUID**—An option to restore the VM BIOS UUID at time of backup, for restore to a new VM as well as VM restore using instant access.
- **Restore individual VMDK**—The ability to restore individual VM disks when restoring back to original or to an alternate VM on the same vCenter or on a different one.

Network considerations

TSDM requires connectivity to PowerProtect DD for data path purposes (see various flows described in Transparent snapshots life cycle). This communication is facilitated using VMkernel (VMK) ports on the ESXi hosts where the TSDM VIB is installed. The transparent snapshots solution would work outside of the box without dedication of VMK ports because any VMK port that can communicate with PowerProtect DD would be automatically used.

That said, for optimal predicted performance and scale, the following guidelines are recommended:

- **Dedicated VMK ports:** Create a single dedicated VMK port per ESXi host. Having a dedicated VMK port decreases the chances of performance degradation due to sharing of VMK ports with other consumers, especially vMotion and vSAN.
- **VMK ports placement:** We recommended placing the VMK port on a VLAN that is dedicated for TSDM to PowerProtect DD traffic or a VLAN dedicated for backup traffic. Having the VMK ports and the relevant PowerProtect DD ports on the same L2 network (same broadcast domain) is advised. Avoid placing the VMK ports and PowerProtect DD ports on VLANs with heavy burst traffic such as vMotion, iSCSI networks, or FT.
- **Consistent end-to-end MTU:** Ensure that the MTU set on the VMK port and on the PowerProtect DD port is uniform from end to end. You can validate this setting by running the ESXCLI command `vmkping` with the DF flag. For example, the following command checks whether there is a uniform end-to-end jumbo frame through a specific VMK port:

```
vmkping -I vmk1 -d -s 8972 10.10.100.1
```

VM Direct Engine considerations The VM Direct Engine (VMDE) runs on the PowerProtect appliance itself and it is available for TSDM operations without having to perform any specific configuration. In terms of scale, TSDM supports up to 180 concurrent VM backups. This embedded VMDE is suitable for all TSDM use cases. TSDM does not use external VMDEs.

Transparent snapshots life cycle

Transparent snapshots provide for a simple, fast, and efficient VM backup. The high-level life cycle is as follows:

- **Monitor:** Track delta changes in memory
- **Process:** Transfer delta changes directly to protection storage
- **Release:** Remove delta table and any temporary data blocks

To provide a better understanding of the VM backup process, this section describes the synchronization and data transfer process, which consists of four major steps:

- Full sync
- Transparent Snapshot creation
- Delta sync
- Snapshot retire

Full sync operation

The full sync operation process, as shown in Figure 6, is as follows:

1. PowerProtect Data Manager issues a full sync request. This request includes all required parameters such as VM information, disk inclusion details, and disk exclusion details.
2. PowerProtect Data Manager queries vCenter to locate the relevant ESXi host, and the operation is transferred to the ESXi host.
3. The ESXi host synchronizes with the TSDM component, leverages VAIO, and makes TSDM aware that a full sync should be performed on the specific asset.
4. The TSDM first uses VAIO to read and query the allocated areas of the disks. After resolving the allocated areas, the TSDM starts to read the data.
5. The TSDM also uses the DD Boost library to establish a connection to the PowerProtect appliance. Empty files are created in the secondary storage (each file corresponding to the flat VMDK file of the VM asset), and the data transfer begins. Eventually, all allocated areas are transferred and written to the PowerProtect appliance.
6. When the full sync operation is complete, the TSDM sends an acknowledgment to the ESXi host. vCenter marks the task as complete.

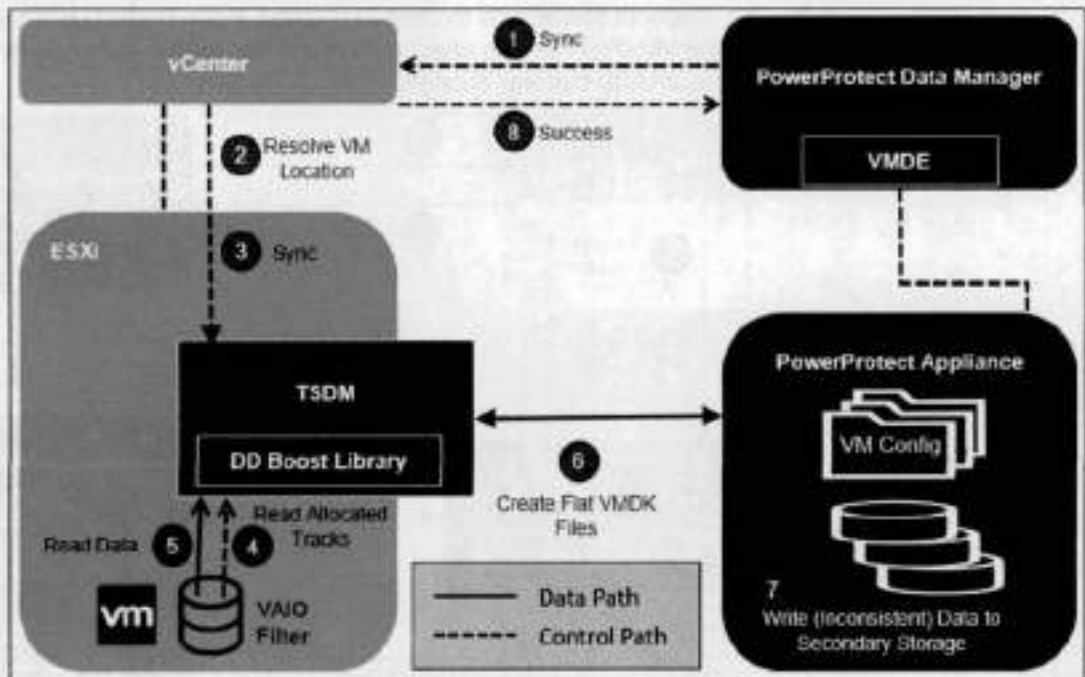


Figure 6. Full sync operation

Note: Because files created during the full sync flow are inconsistent, you cannot use them to restore the VM to a consistent point in time. During the full sync flow, the VM still serves I/O operations, and data in disks might change during the full sync operation itself. A delta sync operation must be performed right after the full sync operation, which creates a consistent point-in-time copy that can be used later for a restore operation.

Transparent snapshots creation

The transparent snapshots creation process (Figure 7), is as follows:

1. After the full sync is complete, PowerProtect Data Manager issues a snapshot creation operation.
2. PowerProtect Data Manager requests a sync operation against vCenter.
3. The API calls are passed to the ESXi host after the location is resolved. In this step, the TSDM has no active role because no data transfer occurs.
4. The ESXi host communicates with the relevant VM asset using VAIO, and the snapshot is persisted to the Snapshot Extent Store (SES).

Note: The Snapshot Extent Store (SES) is dynamically created during snapshot creation and is deleted when the snapshot is retired. The SES stores the bitmaps that correspond to the data in the disk at that time. The SES uses thin-allocated space across the overall datastore space and does not affect the specific VM quota.

5. Using the SES, the VAIO filter takes the bitmap in memory and saves it. Because all of it is bitmap based, creating the transparent snapshot is fast, which reduces the read/write latency.
6. After the bitmap is persisted, the filter can start tracking changes on the disk again, using a new bitmap.

- The snapshot operation is marked as finished. PowerProtect Data Manager can use the VM Direct Engine to access the vCenter level task completion and get the snapshot UUID that was created.

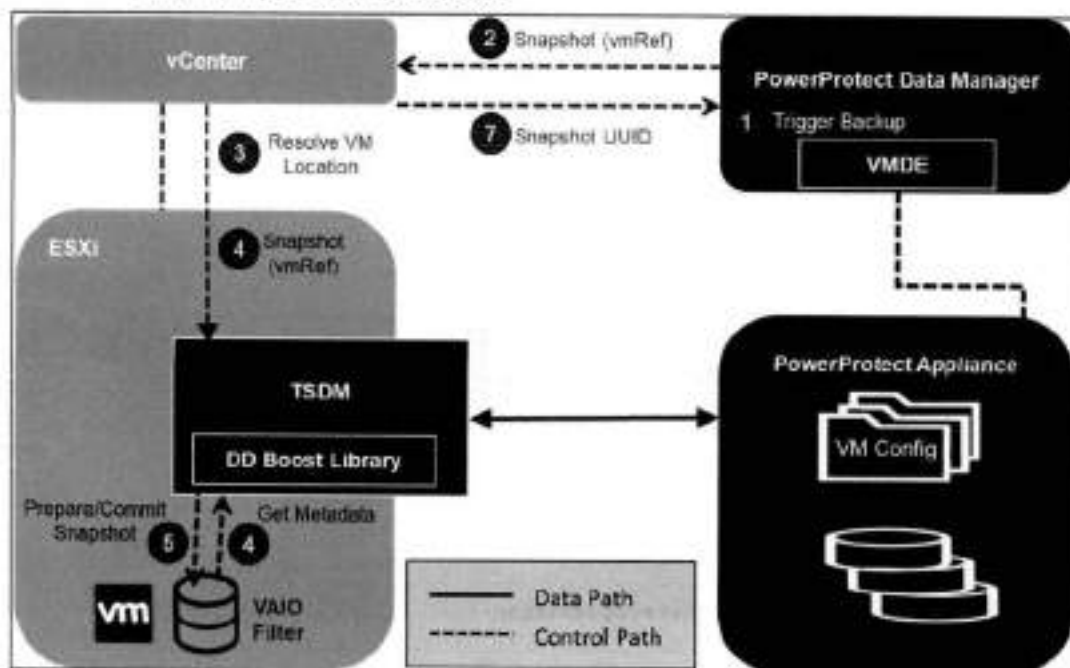


Figure 7. Transparent snapshot creation

Delta sync operation

After the transparent snapshot creation operation is complete, you can initiate a delta sync operation or an incremental operation (see Figure 8). The process is as follows:

- PowerProtect Data Manager issues API calls to vCenter for a delta sync operation and provides the previously created snapshot UUID.
- PowerProtect Data Manager signals to TSDM to start the delta sync flow through vCenter API, which resolves the relevant ESXi host.
- The TSDM uses VAIO APIs to query and track the changed areas that the transparent snapshot bitmap represents.
- For each changed area, the data is read from the disk.

Note: The delta sync operation uses a Fast Copy overwrite approach. In this approach, the previous point-in-time files are first fast copied. The fast-copied files are partly overwritten with the incremental data. Only the delta or changes that are represented by the currently synced snapshot are copied. These changes are copied in the Snapshot Extent Store (SES).

- The changed data or delta is read from the disks to create a consistent data flow.
- The read data is written to protection storage using the DD Boost library.
- The changed data write to protection or secondary storage is now complete.

8. When all data has moved to protection storage, the TSDM sends an acknowledgment to the ESXi host that the operation is complete. The vCenter level task is marked as complete.
9. The metadata is written to protection storage. In this step, the VM metadata (such as VMX files, manifest, and the last TSDM snapshot information) is transferred using the VM Direct Engine VMware APIs.

Note: From this point in time, the files on the PowerProtect appliance are crash consistent and can be used for recovery. For a full backup, both full sync and delta sync are performed. However, for an incremental backup, only delta sync is performed. A full sync can back up four VM disks in parallel, a delta sync can back up 10 VM disks in parallel.

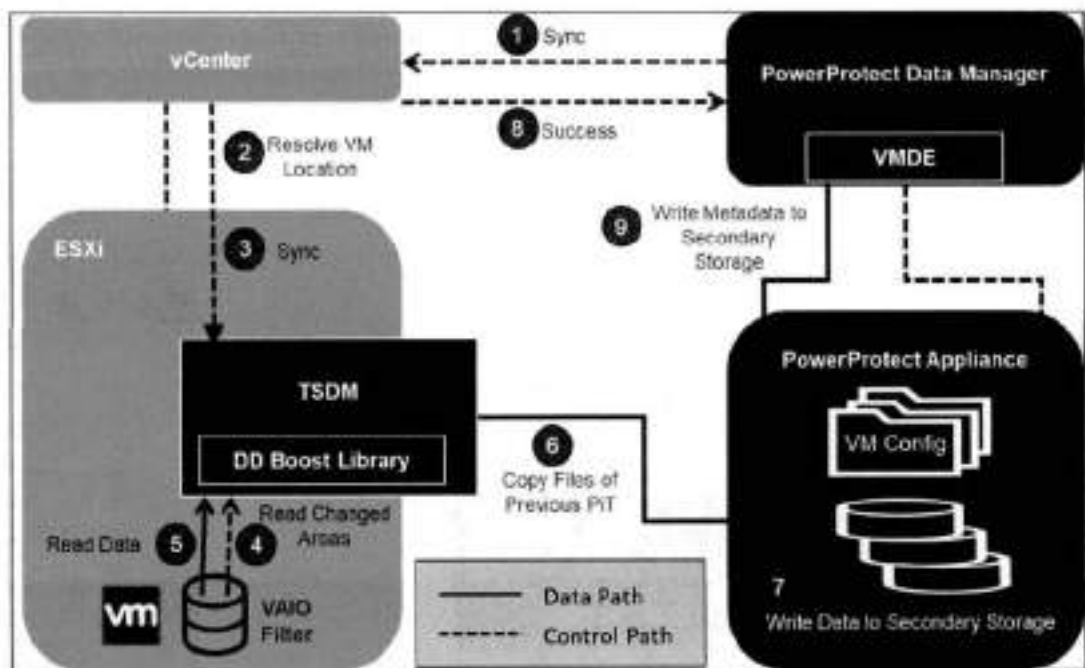


Figure 8. Delta sync operation

Snapshot retire operation

After the delta sync operation is complete, PowerProtect Data Manager ensures that the previously created snapshot is deleted. The snapshot retire operation process (Figure 9) is as follows:

1. PowerProtect Data Manager calls on vCenter to invoke the retire snapshot API towards the relevant ESXi host.
2. The ESXi host relays this information to all the relevant VAIO filters to delete all the bitmaps and copy-on-write data residue left from the snapshot creation and delta sync stages.
3. The ESXi host sends an acknowledgment of the successful snapshot retire operation to vCenter and PowerProtect Data Manager.
4. PowerProtect Data Manager records the backup copy set information in the PowerProtect Data Manager Catalog and informs the search node (if any) to start gathering metadata for indexing.

Restoring virtual machines

Asset	State	Size	Age	Resc.	Summary
DRG-01_North	Success	11.2 GB	30 GB	00:00	Description : Backup for 'DRG0101', total VMs in VMID : 1 SnapshotName : Snapshot TransportMethod : NFS Parallelism : 3 Background : Full Overall : CompressedSize : 424.3 MB Success/Transfer : 31.5 MB/s

Note:

- Transparent snapshots are supported on VMFS, NFS, and vSAN datastores, as well as vVols.
- Virtual and physical RDM volumes are not supported with transparent snapshots.

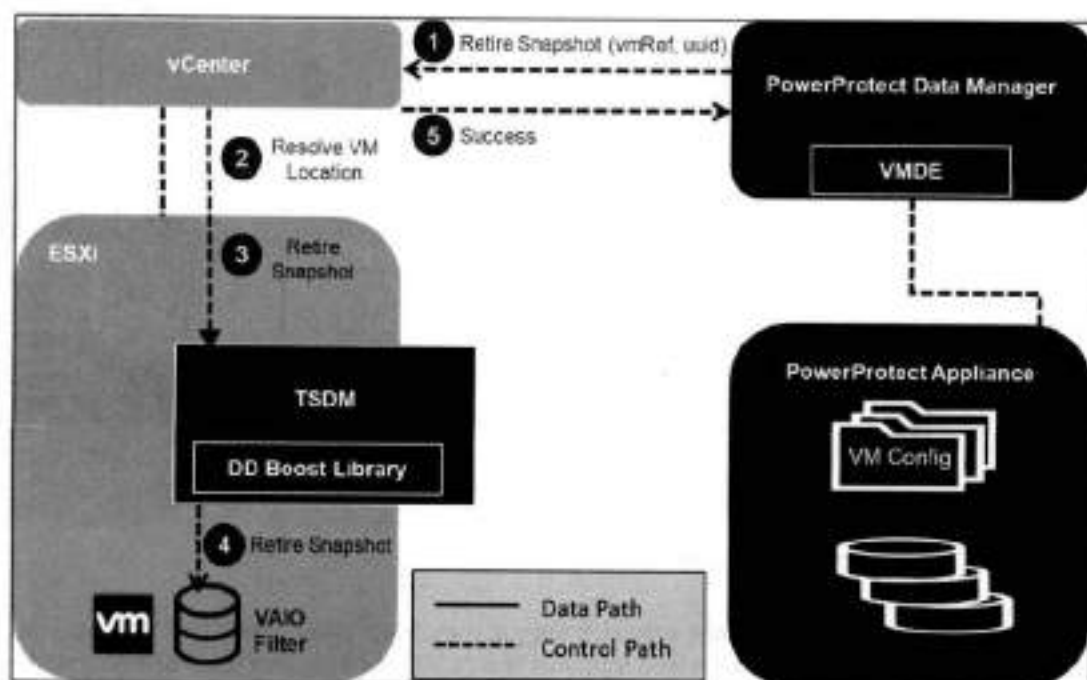


Figure 9. Snapshot retire operation

Restoring virtual machines

Before invoking any APIs from PowerProtect Data Manager to vCenter, PowerProtect Data Manager takes care of the configuration of the VM. For example, if a disk should be added or removed because the VM changed after the snapshot was taken, the disk might be added. For this reason, the reconfiguration part is performed before the virtual machine is restored.

The virtual machine restore process (Figure 10) is as follows:

1. A metadata-only snapshot is taken using the same snapshot creation flow described in Transparent snapshots creation.
2. PowerProtect Data Manager invokes the restore operation, directing which VM should be restored to what point in time.
3. The VM is checked to determine whether it is in a powered off state. If it is not, the VM is powered off.
4. PowerProtect Data Manager locates the relevant ESXi host through vCenter, and the ESXi host communicates with the TSDM to initiate a restore operation.
5. For a restore workflow, first reserve all the areas of the VM disks that should be reverted to the previous point in time, for the following reasons:
 - To minimize the data you transfer from the protection storage back to the disk
 - To identify which parts of each disk have changed since the point in time to which we are trying to revert
6. To resolve the previously described changes, the restore workflow leverages two Get Diff APIs, namely Get VAIO Diff API and DD Get Diff API.
7. The ESXi host resolves the Get VAIO Diff API. This difference includes the changes that the VM has made to the disks since the last point in time that was previously synced or backed up to the PowerProtect appliance. The VAIO Diff uses the metadata only snapshot, taken before the VM was powered off, to get the details on what was never written to the PowerProtect appliance.
8. The DD Get Diff API provides the delta details between the last sync point-in-time and the one to revert. It also merges the delta with the delta returned from the previous step. This step provides the complete set of extents that can now be read from the PowerProtect appliance.
9. From the point-in-time copy to which the user wants to revert, data is read and then finally written on top of the VMDKs.
10. After all the data movement is complete, the TSDM sends an acknowledgment to the ESXi host that the restore process has been completed.
11. vCenter marks the task completed and sends an acknowledgment to PowerProtect Data Manager for catalog update.
12. The VM can now be powered on and should have been successfully reverted to the previous point in time.
13. The metadata only snapshot can now be retired, in the same manner as after every delta sync operation.

Note:

- Multiple streams are opened on the TSDM (when ESXi receives the request) to achieve restore parallelism. You can achieve a higher level of parallelism if the VMs are spread across multiple ESXi hosts.
- Restore in parallel supports up to eight disks of a VM using transparent snapshots.

Performance test results

- Only Restore to Original and New are supported with PowerProtect Data Manager 19.9 release using transparent snapshots. Instant access and File Level Restore do not use a specific data mover; hence, they are supported for TSDM-based backups.

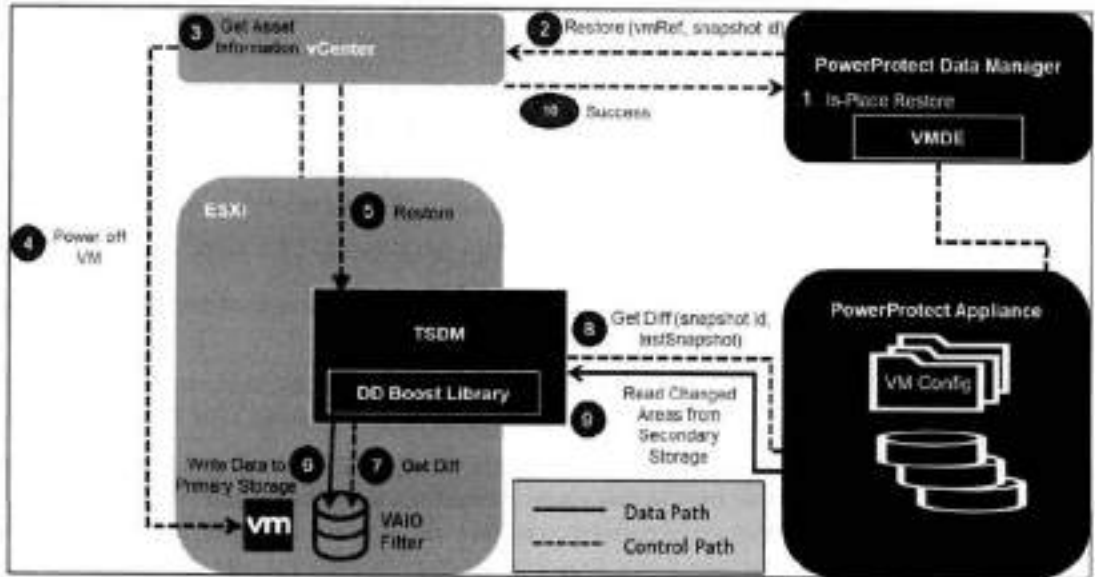


Figure 10. Restoring virtual machines

Performance test results

This section compares the I/O characterization between VMware vSphere Storage APIs - Data Protection (VADP) and transparent snapshots using PowerProtect Data Manager 19.9. You can infer from these results that with transparent snapshots, you overcome the penalties of the write and the read latencies. This result can reduce VM latency by up to five times and provide up to five times faster backups.¹




Figure 11. Backup performance comparison

¹ Disclaimer: These results compare PowerProtect Data Manager 19.9 with transparent snapshots backup performance (performance-optimized mode) to PowerProtect Data Manager with VADP backup performance. The results are based on Dell Technologies internal testing in August 2021.

	Transparent Snapshot	VADP
Test 1	Effective IOPS during Sync: AVR - 10K Sync transfer rate: AVR - 550 MB/s Latency: During Sync: AVR - 1 ms No Sync: AVR - 0.5 ms	Effective IOPS during Sync: AVR - 10K Sync transfer rate: AVR - 110 MB/s Latency: During Sync: AVR - Read: 1 ms / Write: 2.2 ms No Sync: AVR - 0.5 ms
Test 2	Effective IOPS during Sync: AVR - 10K Sync transfer rate: AVR - 550 MB/s Latency: During Sync: AVR - 1 ms No Sync: AVR - 0.5 ms	Effective IOPS during Sync: AVR - 10K Sync transfer rate: AVR - 110 MB/s Latency: During Sync: AVR - Read: 1 ms / Write: 2 ms No Sync: AVR - 0.5 ms
Test 3	Effective IOPS during Sync: AVR - 10K Sync transfer rate: AVR - 610 MB/s Latency: During Sync: AVR - 1 ms No Sync: AVR - 0.5 ms	Effective IOPS during Sync: AVR - 10K Sync transfer rate: AVR - 173 MB/s with 12 DD Streams Latency: During Sync: AVR - Read: 1.1 ms / Write: 2.5 ms No Sync: AVR - 0.5 ms

Backup performance comparison

Figure 12. Backup performance comparison test details

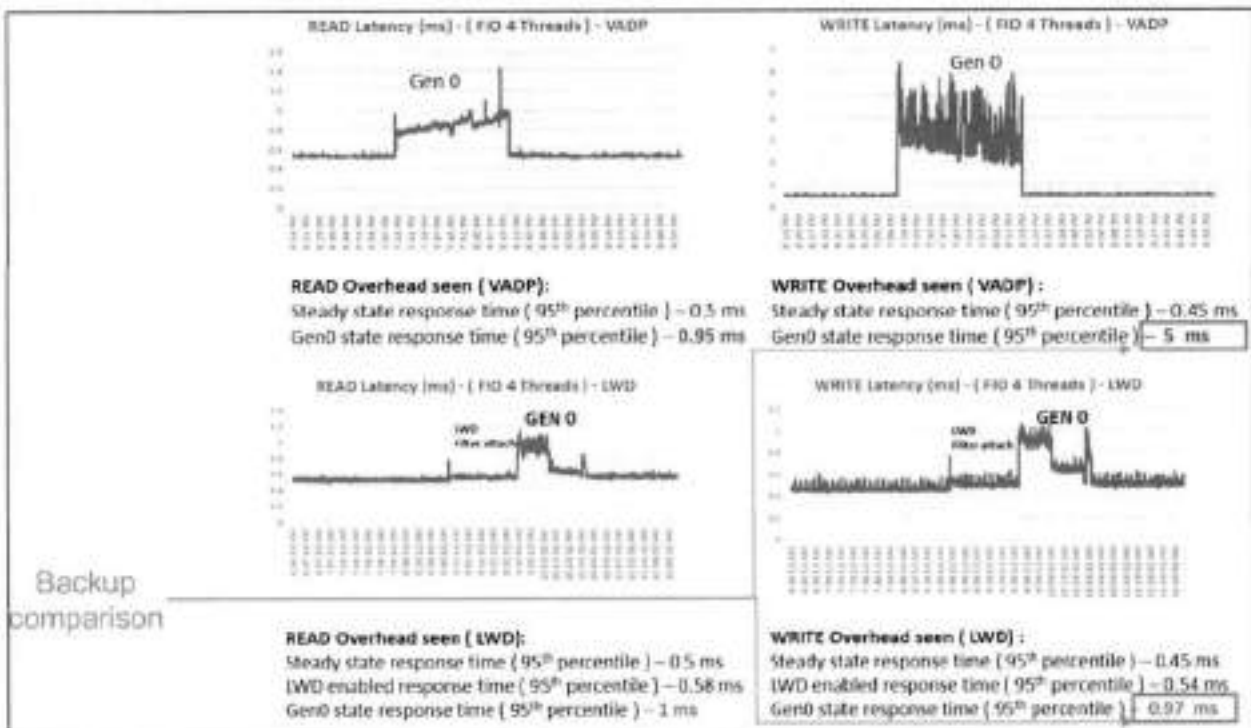


Figure 13. VM latency (read and write) performance comparison

Performance improvements in PowerProtect Data Manager 19.10

Restore performance

The following results show the restore performance improvements of PowerProtect Data Manager 19.10, as compared to PowerProtect Data Manager 19.9.




Figure 14. Restore performance comparison

Test	19.9 Throughput (MB/s)	19.10 Throughput (MB/s)	% Improvement
Restore to New (1 VM Full)	120	298	148.3
Restore to Original (4 VMs Delta -10%)	15171	19924	31.3
Restore to New (4 VMs Full Restore)	329	652	98.2

Figure 15. Restore performance comparison test details

References

Dell Technologies documentation

The following Dell Technologies resources provides other information related to this document. Access depends on your login credentials. If you do not have access to a resource, contact your Dell Technologies representative.

- [PowerProtect Data Manager Interactive Demo](#)
- [PowerProtect Data Manager documentation \(Dell Support\)](#)

VMware documentation

See also the following VMware documentation:

- [VMware vSphere APIs for I/O Filtering \(VAIO\)](#)

Dell EMC PowerProtect Cyber Recovery

Proteção moderna e comprovada para dados críticos contra ransomware e ataques cibernéticos destrutivos

POR QUE ESCOLHER O CYBER RECOVERY?

Os ataques cibernéticos são projetados para destruir, roubar ou comprometer seus dados valiosos, inclusive os backups. Proteger seus dados críticos e recuperá-los com a integridade assegurada é fundamental para retomar as operações de negócios normais após o ataque. Sua empresa conseguiria? Veja a seguir cinco componentes de uma solução de recuperação cibernética moderna e comprovada:

Isolamento e governança de dados

Um ambiente de data center isolado, desconectado de redes corporativas e de backup e restrito a usuários com uma autorização adequada.

Cópia de dados automatizada e air gap

Crie cópias de dados inalteráveis em um cofre digital seguro e processos que criam um air gap operacional entre o ambiente de produção/backup e o cofre.

Ferramentas e lógica analítica inteligentes

Aprendizado de máquina e indexação total de conteúdo com lógica analítica eficiente da segurança do cofre. As verificações de integridade automatizadas determinam se os dados foram afetados por malware, e as ferramentas ajudam na correção, se necessário.

Recuperação e correção

Fluxos de trabalho e ferramentas para realizar a recuperação após um incidente usando processos de restauração dinâmicos e os procedimentos de DR existentes.

Planejamento e design da solução

Orientação especializada para selecionar conjuntos de dados críticos, aplicativos e outros ativos vitais para determinar os RTOs e RPOs, e otimizar a recuperação.

O desafio: ataques cibernéticos são inimigos das empresas orientadas por dados

Os dados são a moeda da economia da internet e um ativo essencial que precisa ser protegido, mantido confidencial e disponibilizado a qualquer momento. O mercado global atualmente se baseia no fluxo constante de dados em redes interconectadas, e as iniciativas de transformação digital colocam mais dados confidenciais em risco.

Isso faz com que os dados da sua organização sejam um alvo atraente e lucrativo para criminosos cibernéticos.


A cibercriminalidade é considerada a maior transferência de riqueza da história, e tudo se baseia nos dados. A Accenture estima que US\$ 5,2 trilhões de valor global estão em risco com a cibercriminalidade nos próximos 5 anos.

Independentemente do setor ou do porte da organização, os ataques cibernéticos expõem continuamente os negócios e os governos a dados comprometidos, perda de receita devido ao tempo de inatividade, danos à reputação e multas regulatórias caras. O custo médio anual de crimes cibernéticos por empresa cresceu para US\$ 13 milhões em 2018, um aumento de 72% nos últimos 5 anos.

Ter uma estratégia de recuperação cibernética se tornou obrigatório para líderes empresariais e governamentais. 79% dos executivos globais classificam ataques cibernéticos como uma das maiores prioridades de gerenciamento de riscos da organização, de acordo com um estudo de 2019 da Marsh & Microsoft.

Então, o que você pode fazer para proteger sua organização e seus dados valiosos?

A solução: PowerProtect Cyber Recovery



Para reduzir os riscos de negócios causados por ataques cibernéticos e criar uma abordagem mais resiliente cibernética à proteção de dados, você pode modernizar e automatizar suas estratégias de recuperação e continuidade dos negócios, além de aproveitar as ferramentas inteligentes mais recentes para detectar e se defender de ameaças cibernéticas.

O Dell EMC PowerProtect Cyber Recovery oferece proteção comprovada, moderna e inteligente para isolar dados essenciais, identificar atividades suspeitas e acelerar a recuperação de dados, permitindo o retorno rápido às operações normais dos negócios.

PowerProtect Cyber Recovery: proteção comprovada, moderna e inteligente para reduzir os riscos de negócios de ameaças cibernéticas

- **Cofre do Cyber Recovery:** o cofre do PowerProtect Cyber Recovery oferece várias camadas de proteção para fornecer resiliência contra ataques cibernéticos mesmo de uma ameaça interna. Ele transfere dados críticos da superfície de ataque, isolando-os fisicamente em uma parte protegida do data center. Para acessar, ele exige credenciais de segurança separadas e autenticação baseada em vários fatores.

As proteções adicionais incluem um air gap operacional automatizado para fornecer isolamento de rede e eliminar as interfaces de gerenciamento que podem ser comprometidas.

O PowerProtect Cyber Recovery automatiza a sincronização de dados entre os sistemas de produção e o cofre, criando cópias imutáveis com políticas de retenção bloqueadas. Se ocorrer um ataque cibernético, você poderá identificar rapidamente uma cópia dos dados, recuperar seus sistemas essenciais e colocar os negócios em funcionamento novamente.



- **CyberSense:** o PowerProtect Cyber Recovery é a primeira solução a integrar totalmente o CyberSense, que adiciona uma camada de proteção inteligente para ajudar a encontrar a corrupção de dados quando um ataque entra no data center. Essa abordagem inovadora fornece indexação total de conteúdo e usa aprendizado de máquina (ML) para analisar mais de cem estatísticas baseadas em conteúdo e detectar sinais de corrupção devido a ransomware. O CyberSense localiza corrupção com até 99,5% de confiança, ajudando na identificação de ameaças e no diagnóstico de vetores de ataque, além de proteger o conteúdo essencial aos negócios, tudo isso na segurança do cofre.
- **Recuperação e correção:** o PowerProtect Cyber Recovery fornece procedimentos automatizados de restauração e recuperação para retomar os sistemas essenciais de negócios de forma rápida e com confiança. Como parte do PowerProtect Data Manager e para os clientes que executam o Dell EMC NetWorker Cyber Recovery, ele permite a recuperação automatizada do cofre. A Dell EMC e seus parceiros de ecossistemas fornecem uma metodologia abrangente para a proteção de dados, bem como a realização de avaliações de danos e a perícia para recuperar seus sistemas ou corrigir e remover o malware ofensivo.
- **Planejamento e design da solução:** os serviços opcionais de consultoria da Dell EMC ajudam a definir quais sistemas essenciais de negócios devem ser protegidos e podem criar mapas de dependência para aplicativos e serviços associados, bem como a infraestrutura necessária para recuperá-los. O serviço também gera requisitos de recuperação e opções de design, além de identificar as tecnologias para analisar, hospedar e proteger seus dados, juntamente com um caso de negócios e implementação da linha do tempo.

Proteger seus dados vitais contra ataques cibernéticos exige soluções comprovadas e modernas. O PowerProtect Cyber Recovery transmite a confiança de que você pode identificar e restaurar rapidamente dados em boas condições e retomar as operações de negócios normais após um ataque cibernético.



[Saiba mais](#) sobre
o Dell EMC
PowerProtect Cyber



[Entre em contato](#) com um
especialista da Dell EMC

[Fonte: estudo "The Cost of Cybercrime Study" da Accenture, 2019]
[Fonte: estudo "The Cost of Cybercrime Study" da Accenture, 2019]
[Fonte: estudo "Global Cyber Risk Perceptions" da Marsh & Microsoft, 2019]

© 2020 Dell Inc. ou suas subsidiárias. Todos os direitos reservados. Dell, EMC e outras marcas comerciais são de propriedade da Dell Inc. ou de suas subsidiárias. Outras marcas comerciais podem pertencer a seus respectivos proprietários. Número de referência: H17020

PowerProtect Data Manager 19.14 Compatibility

September 13, 2023

Contents

DEPLOYMENT.....	2
BACKUP TARGET.....	3
LANGUAGE SUPPORT.....	4
VMWARE.....	5
SQL.....	8
EXCHANGE.....	8
ORACLE.....	9
* HANA.....	10
STORAGE ARRAYS VIA REST API.....	10
STORAGE DATA MANAGEMENT.....	11
DYNAMIC NAS.....	11
FILE SYSTEMS.....	12
POWERPROTECT DATA MANAGER IN CLOUD MARKETPLACES.....	14
CLOUD DISASTER RECOVERY.....	15
KUBERNETES.....	17
SUPPORTASSIST.....	18
SYSLOG FORWARDING.....	18

DEPLOYMENT

Platform	Platform Version	Comments
VMware vCenter/ESXi	6.7 U3, 7.0, 7.0 U1, 7.0 U2, 7.0 U3, 8.0, 8.0 U1	
Amazon Web Services (AWS)	AWS instance type: m5-2xlarge (8 vCPUs, 32GB memory)	Subscribe via AWS Marketplace. Delivered via CloudFormation template, which also optionally deploys PowerProtect DD Virtual Edition 7.11.0.0.
Azure	Azure VM size: Standard D6s v3 (8 vCPUs, 32GB memory)	Subscribe via Azure Marketplace. Delivered via Azure Resource Manager (ARM) template, which also optionally deploys PowerProtect DD Virtual Edition 7.11.0.0.
Google Cloud Platform (GCP)	GCP VM size: Customized (8 vCPUs, 32GB memory)	Subscribe via GCP Marketplace. Delivered via Deployment Manager (DM) template, which also optionally deploys PowerProtect DD Virtual Edition 7.11.0.0.
VMware Cloud (VMC) on AWS	SDDC 1.20, 1.22	Supported when deployed to a VMC on AWS vCenter or to the AWS native cloud.
VMware Cloud on Dell	SDDC 1.20, 1.22	Supported when deployed to a VMC on Dell vCenter.
Azure VMware Solutions (AVS)		For deployments in AVS, software is available from Dell Technologies (not the Marketplace). Subscribe via Azure Marketplace for DDVE deployment. Supported when deployed to an AVS vCenter or to the Azure native cloud.
Google Cloud VMware Engine (GCVE)		For deployments in GCVE, software is available from Dell Technologies (not the Marketplace). Subscribe via GCP Marketplace for DDVE deployment. Supported when deployed to a GCVE vCenter or to the GCP native cloud.
Oracle Cloud VMware Solution (OCVS)		For deployments in OCVS, software is available from Dell Technologies (not the Marketplace).
Web browser	Google Chrome	The latest version of the Google Chrome browser to access the PowerProtect Data Manager UI.

Notes:

- Supported datastores: vVols, vSAN, VMFS, and NFS.
- Includes VMware Cloud Foundation (VCF) versions that contain supported vCenter/vSphere versions.
- When deploying to a VMC on AWS, VMC on Dell, AVS, or GCVE vCenter, deploy PowerProtect Data Manager from the OVA and use the deployment wizard to select the corresponding configuration for this VMC environment.
- PowerProtect Data Manager only supports in-cloud VM image protection in the respective environments. Protection from in-cloud to on-premises environments, or from on-premises to in-cloud environments, is not supported.

BACKUP TARGET

Data Domain (Physical & Virtual)

DDOS	6.2.1, 7.0, 7.1.0, 7.2.0, 7.3.0, 7.4.0, 7.5.0, 7.6.0, 7.7.x, 7.8.0, 7.9.0, 7.10.x, 7.11.0, 7.12.0	
DDMC	6.2.1, 7.0, 7.1.0, 7.2.0, 7.3.0, 7.4.0, 7.5.0, 7.6.0, 7.7.x, 7.8.0, 7.9.0, 7.10.x, 7.11.0, 7.12.0	
DDVE	4.0, 7.0, 7.1.0, 7.2.0, 7.3.0, 7.4.0, 7.5.0, 7.6.0, 7.7.x, 7.8.0, 7.9.0, 7.10.x, 7.11.0, 7.12.0	DDVE 4.0 corresponds to DDOS 6.2.

Storage Units

PowerProtect DD System Model	Maximum Number of Storage Units	Configurable Concurrently Active Storage Units
9800, 9900	256	256
6800, 6900, 9300, 9400	128	128
All other supported DD systems	100	32, based on the model
PowerProtect DD Virtual Edition (DDVE)		
4, 6, 8 TB	100	6
12, 48 TB	100	14
64, 96 TB	100	32

Cloud Tier

Vendor Name	Cloud Service
Dell Technologies	ECS (Elastic Cloud Storage)
Microsoft	Microsoft Azure
Amazon	AWS (Amazon Web Services)
Google	GCP (Google Cloud Platform)

Notes:

- DDBoost library packaged in PowerProtect Data Manager OVA: 7.12.0.0-1052569

LANGUAGE SUPPORT

Supported Asset Languages

	Operating System Languages	Notes
Full Support	English	
Conditional Support	Simplified Chinese French German Italian ¹ Japanese	Protection of data in these language environments must observe the following considerations: <ul style="list-style-type: none">• Application agents must be installed to the default location.• Names of assets such as files, folders, and databases can use characters in any supported language.• Hostnames, fully qualified domain names, and cluster names can only use the ASCII characters A-Z, a-z, 0-9, and hyphen.• The operating systems of the backup source and restore destination must use the same locale.

Notes:

- PowerProtect Data Manager and its user interface support only the ASCII characters A-Z, a-z, 0-9, and hyphen in the following functions:
 - APIs
 - Character string input and output.
 - Configuration settings.
 - Exported logs.
 - Headings and fields.
 - Informational and error messages.
 - Local time and date formats.
 - Sorting, searching, and filtering.

1. For VMware virtual machines, Italian language support requires vSphere 8.0 and above.

VMWARE

Platform

Platform	Platform Version	Comments
vCenter	6.7 U3, 7.0, 7.0 U1, 7.0 U2, 7.0 U3, 8.0, 8.0 U1	Includes patch releases, for example: 7.0 U1d, 7.0 U3c.
VMC on AWS	SDDC 1.20, 1.22	Includes odd numbered optional releases.
VMC on Dell	SDDC 1.20, 1.22	Includes odd numbered optional releases.
Azure VMware Solutions (AVS)		
Google Cloud VMware Engine (GCVE)		
Oracle Cloud VMware Solutions (OCVS)		

Notes:

- Supported datastores: vVol, vSAN, VMFS, NFS.
- Supported vSAN topologies: Standard Cluster, Stretched Cluster, Two-node Cluster, HCI Mesh.
- vSAN Express Storage Architecture (ESA) is currently not supported.
- PowerProtect Data Manager supports in-cloud VM image protection in the respective environments. Protection from in-cloud to on-premises environments, or from on-premises to in-cloud environments, is not supported.
- VMware has declared End of General Support for vCenter 6.7. VMware fixes for these products are no longer provided (<https://lifecycle.vmware.com/>).

Hypervisor

Platform	Platform Version	Comments
ESXi	6.7 U3, 7.0, 7.0 U1, 7.0 U2, 7.0 U3, 8.0, 8.0 U1	Includes patch releases, for example: 7.0u1d, 7.0u3c.

Notes:

- VMware has declared End of General Support for ESXi 6.5 and 6.7. VMware fixes for these products are no longer provided (<https://lifecycle.vmware.com/>).

Management (vSphere Plugin)

Platform	Platform Version	Comments
vCenter	6.7 U3, 7.0, 7.0 U1, 7.0 U2, 7.0 U3, 8.0, 8.0 U1	7.0 supported with Classic View.

Virtual Machines

Platform	Platform Version	Comments
VMware Tools	10.x, 11.x, 12.x	Version 11.1 and later is required to perform Microsoft SQL Server application-aware protection.
Open VM Tools	10.x, 11.x, 12.x	For Linux platforms.

Notes:

- All virtual machine operating systems supported by VMware.

Transparent Snapshot Data Mover (TSDM)

Supported Features & Products	Platform Version	Comments
VMware Virtual Volumes (vVols)	vSphere 7.0u3d (and above)	
CyberSense		https://www.indexengines.com/cs
VMware Cloud Foundation (VCF)	4.4 (and above)	Bill of Materials (BOM) contains vSphere 7.0 U3c
Oracle Cloud VMware Solution (OCVS)	vSphere 7.0 U3d (and above)	
Microsoft SQL Application Aware Protection	vSphere 7.0 U3c (and above)	VM Application Aware Backup – SQL provides supported versions.


vRealize Data Protection Extension

Platform	Platform Version	Comments
VMware vRealize Automation with Embedded vRealize Orchestrator	8.8.x, 8.9.x, 8.10.x, 8.11.x	Includes standalone and cluster deployments.
VMware Aria Automation with Embedded Aria Orchestrator	8.12.x	Includes standalone and cluster deployments.

VCF (on VxRail) Management Protection

Platform	Platform Version	Comments
VMware Cloud Foundation	4.4.x, 4.5.x	

Notes:

- Includes VCF on vSAN Ready Nodes and VCF on VxRail.
- Configure protection for VCF management components with the `ppdm-vof-component-protection.sh` script that you obtain through the UI by clicking  > **Downloads > VMware** and then selecting **VMware Cloud Formation (VCF) on VxRail**.

VMware Telco Cloud Platform

Platform	Platform Version	Comments
Telco Cloud Platform (TCP)	2.5	VMware Tanzu Kubernetes Grid (TKG) 1.5.x, ESXi and vCenter 7.0 U3 (including patch releases)
	2.7	VMware Tanzu Kubernetes Grid 1.6.x, ESXi and vCenter 7.0 U3 (including patch releases)
	3.0	VMware Tanzu Kubernetes Grid 2.1.x, ESXi and vCenter 8.0 (including patch releases)

Notes:

- For TCP components that require file-level protection, refer to the *File Systems* table on page 12 for supported file and operating systems.
- For TKGm workload clusters, refer to the *Kubernetes* table on page 17 for supported distributions.
- For more information, refer to the *Dell Technologies PowerProtect Data Manager for VMware Telco Cloud Platform Solution Brief* and the *VMware Telco Cloud Platform Release Notes*.

VM Application Aware Backup – SQL

Platform	Platform Version	Comments
VMware vSphere (ESXi and vCenter)	6.7 U3, 7.0, 7.0 U1, 7.0 U2, 7.0 U3, 8.0, 8.0 U1	SQL TSDM supports VMware vSphere 7.0 U3c (and above)
Microsoft SQL Server Version	MS SQL Server Bitness	Windows OS Version [Hardware architecture-64 bit]
2014 Service Pack 3 (supported up to CU4)	32-bit, 64-bit	2012, 2012 R2, 2016
2016 Service Pack 3	64-bit	2012, 2012 R2, 2016, 2019
2017 (supported up to CU31)	64-bit	2012, 2012 R2, 2016, 2019, 2022
2019 (supported up to CU20)	64-bit	2016, 2019, 2022
2022 CU4 and above only	64-bit	2016, 2019, 2022

Notes:

- Clustered and cluster-less AAG Quorum supports file share witness only.
- VMs with multi-writer/shared disks are not supported.
- VMware Tools: 11.x and later. Clustered AAG for pure IPv6 requires VMware Tools 12.x.

VM File Level Recovery

Operating System	Hardware Architecture	OS Vendor	Filesystems	Search and Indexing
Windows 10	64-bit	Microsoft	NTFS	Yes
Windows 11	64-bit	Microsoft	NTFS	Yes
Windows 2012	64-bit	Microsoft	NTFS, FAT32, ReFS	Yes ²
Windows 2012 R2 ¹	64-bit	Microsoft	NTFS	Yes
Windows 2016	64-bit	Microsoft	NTFS	Yes
Windows 2019	64-bit	Microsoft	NTFS, FAT32, Dynamic	Yes
Windows 2022	64-bit	Microsoft	NTFS, ReFS, Dynamic	Yes ²
CentOS 7.1, 7.4, 7.6, 7.8	X64	CentOS	XFS	Yes
CentOS 7.9	X64	CentOS	XFS, BTRFS ³ , EXT4	Yes
CentOS 8 Stream	X64	CentOS	XFS	Yes
CentOS 9 Stream	X64	CentOS	XFS, EXT4	Yes
Red Hat Enterprise Linux (RHEL) 7.4 – 7.8	X64	Red Hat	XFS, EXT2/3/4	Yes
RHEL 7.9	X64	Red Hat	XFS, EXT4, BTRFS ³	Yes
RHEL 8.1 – 8.5	X64	Red Hat	XFS, EXT2/3/4	Yes
RHEL 8.6 – 8.8	X64	Red Hat	XFS, EXT4	Yes
RHEL 9 – 9.2	X64	Red Hat	XFS, EXT4	Yes
SUSE Linux Enterprise Server (SLES) 12 SP5	X64	SUSE	XFS, EXT4	Yes
SLES 15 SP3	X64	SUSE	XFS	Yes
SLES 15 SP4	X64	SUSE	XFS, EXT4, BTRFS ³	Yes
Ubuntu 14.04, 16.04, 18.04, 20.04 Server	X64	Canonical	EXT4, XFS	Yes
Ubuntu 22.04 Server & Desktop	X64	Canonical	EXT4, XFS	Yes
Debian 10.13	X64	Debian	EXT4	Yes
Debian 11.3 – 11.7	X64	Debian	EXT4	Yes
Debian 12.0	X64	Debian	EXT4	Yes
Oracle Linux 7.7	X64	Oracle	XFS, EXT4	Yes
Oracle Linux 7.9	X64	Oracle	XFS, EXT4, BTRFS ³	Yes
Oracle Linux 8.3 – 8.4	X64	Oracle	XFS, BTRFS ³	Yes
Oracle Linux 8.5 – 8.8	X64	Oracle	XFS, EXT4, BTRFS ³	Yes
Oracle Linux 9.0 – 9.2	X64	Oracle	XFS, EXT4	Yes
OpenSUSE Leap 15.4	X64	openSUSE	EXT4, BTRFS ³	Yes
Fedora Linux 37 – 38	X64	Fedora	XFS	Yes
Rocky Linux 8.6 – 8.8, 9.0 – 9.1	X64	Rocky	EXT4, XFS	-

Notes

- Indexing of LVM thin-pool logical volumes is not supported.
 - PowerProtect Data Manager does not support file-level restore to a target VM that is protected using Site Recovery Manager (SRM) or vSphere Replication (VR).
 - The target VM for the restore should support the same filesystem type, version, and options as the source backup. For example, the versions of xFoprogs on the target and source VMs must be compatible.
2. Includes Extended Security Update (ESU) from Microsoft (<https://docs.microsoft.com/en-US/lifecycle/fag/extended-security-updates>).
 3. Indexing of ReFS is not supported.
 4. Restore of files from BTRFS sub-volumes is not supported, only image-level recovery.

SQL

Microsoft SQL Server Version	Server Bitness	Microsoft App Agent Version	OS	Hardware Architecture	OS Vendor	OS Vendor Version
2014 SP3 Supported up to CU4	32-bit, 64-bit	19.10, 19.11, 19.12, 19.13, 19.14	Windows	X64	Microsoft	2012, 2012 R2, 2016, 2019
2016 SP3 (along with latest CUs)	64-bit	19.10, 19.11, 19.12, 19.13, 19.14	Windows	X64	Microsoft	2012, 2012 R2, 2016, 2019
SQL Server 2017 CU31 (latest)	64-bit	19.10, 19.11, 19.12, 19.13, 19.14	Windows	X64	Microsoft	2012, 2012 R2, 2016, 2019, 2022
SQL Server 2019 CU20 (latest)	64-bit	19.10, 19.11, 19.12, 19.13, 19.14	Windows	X64	Microsoft	2016, 2019, 2022
SQL Server 2022 CU 3	64-bit	19.10, 19.11, 19.12, 19.13, 19.14	Windows	X64	Microsoft	2016, 2019, 2022

Notes:

- Application Agent for MS SQL supports:
 - MS SQL clustering for AAG and FCI configurations.
 - MS SQL clusterless AAG configurations.
 - All combinations of OS and application versions in this table, provided that the OS is supported by the application.
 - All file systems shown in the table, provided that these file systems are supported by the application.
 - All Windows Server editions and all Service Packs (SP) for a given release, provided that the combination is supported by the application.
 - All SQL server editions such as Standard, Enterprise, Express, Workgroup, and Web that are supported by the application release.
 - SSMS 2012 through 19.x versions in SQL Servers based on Microsoft compatibility.

EXCHANGE

Microsoft Exchange Version	Server Bitness	Microsoft App Agent Version	OS	Hardware Architecture	OS Vendor	OS Vendor Version
2013	64-bit	19.5, 19.6, 19.7, 19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Windows	X64	Microsoft	2012, 2012 R2
2016	64-bit	19.5, 19.6, 19.7, 19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Windows	X64	Microsoft	2012, 2012 R2, 2016
2019	64-bit	19.5, 19.6, 19.7, 19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Windows	X64	Microsoft	2019, 2022

Notes:

- Application Agent for Exchange supports:
 - All combinations of OS and application versions in the table, provided that the OS is supported by the application.
 - All virtualization environments (such as VMware) at the guest level, provided that the application and database app agent running on this guest OS natively support this OS.
 - All Cumulative Updates (CU) and Update Rollups (UR) for the application versions listed in this table unless explicitly stated otherwise.
- PowerProtect Data Manager for Exchange 19.13 and later supports Exchange 2019 on Windows Server Core Edition 2019 and 2022.
- PowerProtect Data Manager for Exchange only supports Windows Server 2022 from Exchange 2019 CU11 and above.
- Because Outlook cannot be installed on Windows Server Core Editions, ItemPoint is not supported directly on Windows Server Core Edition. Instead, use the proxy server method for ItemPoint operations.
- Co-existence of the Exchange agent is supported only with the File System agent.

ORACLE

Oracle Version	Oracle Bitness	RMAN App Agent Version	OS	Hardware Architecture	OS Distribution	OS Vendor Version
11g R2	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	Red Hat Enterprise Linux	7
11g R2	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	Oracle Linux	7
11g R2	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	SLES	12
11g R2	64-bit	19.9, 19.10, 19.11, 19.12, 19.13, 19.14	IBM AIX	X64	AIX	7.1, 7.2
12c R1	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	Red Hat Enterprise Linux	7
12c R1	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	Oracle Linux	7
12c R1	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	SLES	12
12c R1	64-bit	19.9, 19.10, 19.11, 19.12, 19.13, 19.14	IBM AIX	X64	AIX	7.1, 7.2
12c R2	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	Red Hat Enterprise Linux	7
12c R2	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	Oracle Linux	7
12c R2	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	SLES	12
12c R2	64-bit	19.9, 19.10, 19.11, 19.12, 19.13, 19.14	IBM AIX	X64	AIX	7.1, 7.2
12c	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	Red Hat Enterprise Linux	7
18c	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	Oracle Linux	7
18c	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	SLES	12
18c	64-bit	19.9, 19.10, 19.11, 19.12, 19.13, 19.14	IBM AIX	X64	AIX	7.1, 7.2
19c	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	Red Hat Enterprise Linux	7, 8
19c	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	Oracle Linux	7, 8
19c	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	SLES	12, 15
19c	64-bit	19.9, 19.10, 19.11, 19.12, 19.13, 19.14	IBM AIX	X64	AIX	7.1, 7.2, 7.3
21c	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	SLES	15
21c	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	Oracle Linux	7, 8
21c	64-bit	19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	Red Hat Enterprise Linux	8

- Notes:
- PowerProtect Data Manager for Oracle supports:
 - Oracle Server when installed and configured on one of the following systems:
 - Stand-alone configuration.
 - Oracle RAC in active/active configuration only.
 - Data Guard configuration.
 - Oracle RAC on all Cluster File Systems supported by Oracle.
 - All combinations of OS and application versions in the table, provided that the OS is supported by the application.
 - All Fix Packs/Patch Sets for all application versions in this table.
 - All Service Packs, Updates, Patches, etc. of all OSs listed in this table, provided that the update is supported by the application (for example, stated support for Oracle Linux 7 also implies support for Oracle Linux 7 Update 5).
 - Raw devices and all file systems shown in the table, provided that these file systems are supported by the application.
 - All virtualization environments (such as VMware, LPAR) at the guest level, provided that the application and database app agent running on this guest OS natively support this OS.
 - Oracle Incremental Merge (OIM), but only on the Linux platform.
 - Oracle database single-node appliance X7 and X8 models.
 - PowerProtect Data Manager for Oracle on AIX supports AIX 7.1 TL5 SP4 or later, and AIX 7.2 TL2 SP1 or later.
 - PowerProtect Data Manager for Oracle does not support:
 - An Oracle database installation in a High Availability (active/passive) OS and RAC-based cluster configuration.

SAP HANA

SAP HANA Version	SAP HANA Bitness	DB App Agent Version	OS	Hardware Architecture	OS Distribution	OS Vendor Version
1.x	64-bit	19.8, 19.9, 19.10, 19.11, 19.12	Linux	X64	Red Hat Enterprise Linux	6, 7, 8
1.x	64-bit	19.8, 19.9, 19.10, 19.11, 19.12	Linux	X64	SLES	11, 12, 15
2.x	64-bit	19.9, 19.10, 19.11, 19.12, 19.13	Linux	X64	Red Hat Enterprise Linux	6
2.x	64-bit	19.9, 19.10, 19.11, 19.12, 19.13	Linux	X64	SLES	11
2.x	64-bit	19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	Red Hat Enterprise Linux	7, 8, 9
2.x	64-bit	19.10, 19.11, 19.12, 19.13, 19.14	Linux	X64	SLES	12, 15

Notes:

- Application Agent for SAP Hana supports:
 - All combinations of OS and application versions in the table, provided that the OS is supported by the application.
 - All Fix Packs/Patch Sets for all application versions in this table.
 - All Service Packs, Updates, Patches, etc. of all OSs listed in this table, provided that the update is supported by the application.
 - Raw devices and all file systems shown in the table, provided that these file systems are supported by the application.
 - All virtualization environments (such as VMware) at the guest level, provided that the application and database app agent running on this guest OS natively support this OS.

STORAGE ARRAYS VIA REST API

Storage Array Model	Storage Array Operating System Version	Vendor
PowerStore 500T, 1000T, 1200T, 3000T, 3200T, 5000T, 5200T, 7000T, 9200T	3.5.0.0-2050321-retail	Dell

Notes:

- This table applies for protection of block volume and volume group data. For protection of PowerStore NAS shares, refer to the Dynamic NAS table on page 11.
- PowerStore integration does not depend on the Storage Direct Agent service. PowerProtect Data Manager communicates directly with PowerStore cluster over a REST API interface.
- The PowerStore appliance communicates with PowerProtect DD over IP (DD Boost). Storage array protection supports physical PowerProtect DD appliances, as well as DDVE on-premises and in a public AWS cloud.
- Block volume protection policies do not support Metro volumes.
- If a volume or volume group is already part of a PowerStore Manager protection policy, the volume or volume group cannot be protected by PowerProtect Data Manager, and conversely.
- A PowerProtect DD storage unit that is configured as a remote system for a PowerStore Manager protection policy cannot be shared with a PowerProtect Data Manager protection policy, and conversely.
- There is no specific hard limit on the number of PowerStore clusters or appliances that can be managed by a single instance of PowerProtect Data Manager. However, other PowerProtect Data Manager limitations may impose such limits.

STORAGE DATA MANAGEMENT

Storage Array Model	Storage Array Operating System Version	Host Operating System	Hardware Architecture	Storage Direct Agent Version
VMAX3 Hybrid	5977.691.684	Windows 2012, 2012 R2, 2016 Red Hat Enterprise Linux 6.x, 7.x SLES 11, 12, 15	X64	19.2, 19.3, 19.4, 19.5, 19.6, 19.7, 19.8, 19.9, 19.10, 19.11, 19.12, 19.13
	5977.691.684	Windows 2012, 2012 R2, 2016 Red Hat Enterprise Linux 7.x SLES 12, 15	X64	19.2, 19.3, 19.4, 19.5, 19.6, 19.7, 19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14
VMAX All Flash	5977.691.684 (Service Pack ID 5560) and up	Windows 2012, 2012 R2, 2016 Red Hat Enterprise Linux 6.x, 7.x SLES 11, 12, 15	X64	19.2, 19.3, 19.4, 19.5, 19.6, 19.7, 19.8, 19.9, 19.10, 19.11, 19.12, 19.13
	5977.691.684 (Service Pack ID 5560) and up	Windows 2012, 2012 R2, 2016 Red Hat Enterprise Linux 7.x SLES 12, 15	X64	19.2, 19.3, 19.4, 19.5, 19.6, 19.7, 19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14
PowerMax 100/8000)	5978.221.221, 5978.479.479, 5978.669.669, 5978.711.711	Windows 2012, 2012 R2, 2016 Red Hat Enterprise Linux 6.x, 7.x SLES 11, 12, 15	X64	19.2, 19.3, 19.4, 19.5, 19.6, 19.7, 19.8, 19.9, 19.10, 19.11, 19.12, 19.13
	5978.221.221, 5978.479.479, 5978.669.669, 5978.711.711	Windows 2012, 2012 R2, 2016 Red Hat Enterprise Linux 7.x SLES 12, 15	X64	19.2, 19.3, 19.4, 19.5, 19.6, 19.7, 19.8, 19.9, 19.10, 19.11, 19.12, 19.13, 19.14

- Notes:
- Storage Direct supports all service packs, updates, patches, etc. of all OSs listed in this table.
 - Storage Direct supports all virtualization environments (such as VMware) at the guest level for both Windows and Linux (RHEL/SLES).
 - Solutions Enabler 8.4.x and 9.0.x are supported.
 - SMIS server version 8.4 and above is supported.
 - Storage Data Management supports only DDR. There is no support for DDVE.
 - Storage Data Management does not support PowerMax 2500/8500 models.
 - PowerProtect Data Manager continues to support SMIS servers as an asset source for Storage Direct. However, support is subject to change in future releases. It is recommended to use the stand-alone Storage Direct agent with new deployments that require long-term support.

DYNAMIC NAS

Storage Array	Storage Vendor	Supported Protocols
Unity 4.5, 5.0.x, 5.1.x, 5.2.x	Dell	Share options: CIFSv3, NFSv3, NFSv4
PowerStore 1.0.0.0, 1.0.4.0, 2.0.0.0, 2.1.1.0, 3.0, 3.2, 3.5	Dell	Share options: CIFSv3, NFSv3, NFSv4
PowerScale OneFS 8.2.x, 9.0, 9.1, 9.2, 9.2.1, 9.3, 9.4.0.13	Dell	Share options: CIFSv3, NFSv3, NFSv4
Generic NAS protection workflow	*	Share options: CIFSv3, NFSv3, NFSv4

- Notes:
- The generic NAS protection workflow involves backing up a live mount of the target share, without taking a snapshot.
 - The generic NAS protection workflow supports all NAS appliances and file servers.
 - The release schedule of each storage array in this table may differ from that of PowerProtect Data Manager. Contact [Customer Support](#) before you deploy a storage array release that is not listed in this table.

FILE SYSTEMS

Operating System	Hardware Architecture	OS Vendor	Feature Options
Windows 2022	64-bit	Microsoft	Filesystem: NTFS, ReFS
Windows 2012/2012R2	64-bit	Microsoft	Filesystem: NTFS, ReFS
Windows 2016	64-bit	Microsoft	Filesystem: NTFS, ReFS
Windows 2019	64-bit	Microsoft	Filesystem: NTFS, ReFS
Windows 10	64-bit	Microsoft	Filesystem: NTFS
Windows 11	64-bit	Microsoft	Filesystem: NTFS
Red Hat Enterprise Linux 7.0	X64	Red Hat	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 3.10.0-123
Red Hat Enterprise Linux 7.1	X64	Red Hat	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 3.10.0-229
Red Hat Enterprise Linux 7.2	X64	Red Hat	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 3.10.0-327
Red Hat Enterprise Linux 7.3	X64	Red Hat	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 3.10.0-514
Red Hat Enterprise Linux 7.4	X64	Red Hat	Filesystem: ext3, ext4, xfs. Kernel version: 3.10.0-693
Red Hat Enterprise Linux 7.5	X64	Red Hat	Filesystem: ext3, ext4, xfs. Kernel version: 3.10.0-862
Red Hat Enterprise Linux 7.6	X64	Red Hat	Filesystem: ext3, ext4, xfs. Kernel version: 3.10.0-957
Red Hat Enterprise Linux 7.7	X64	Red Hat	Filesystem: ext3, ext4, xfs. Kernel version: 3.10.0-1062
Red Hat Enterprise Linux 7.8	X64	Red Hat	Filesystem: ext3, ext4, xfs. Kernel version: 3.10.0-1127
Red Hat Enterprise Linux 7.9	X64	Red Hat	Filesystem: ext3, ext4, xfs. Kernel version: 3.10.0-1160
Red Hat Enterprise Linux 8.0	X64	Red Hat	Filesystem: ext3, ext4, xfs. Kernel version: 4.18.0-80
Red Hat Enterprise Linux 8.1	X64	Red Hat	Filesystem: ext3, ext4, xfs. Kernel version: 4.18.0-147
Red Hat Enterprise Linux 8.2	X64	Red Hat	Filesystem: ext3, ext4, xfs. Kernel version: 4.18.0-193
Red Hat Enterprise Linux 8.3	X64	Red Hat	Filesystem: ext3, ext4, xfs. Kernel version: 4.18.0-240
Red Hat Enterprise Linux 8.4	X64	Red Hat	Filesystem: ext3, ext4, xfs. Kernel version: 4.18.0-305
Red Hat Enterprise Linux 8.5	X64	Red Hat	Filesystem: ext3, ext4, xfs. Kernel version: 4.18.0-348
Red Hat Enterprise Linux 8.6	X64	Red Hat	Filesystem: ext3, ext4, xfs. Kernel version: 4.18.0-372
Red Hat Enterprise Linux 8.7	X64	Red Hat	Filesystem: ext3, ext4, xfs. Kernel version: 4.18.0-425
Red Hat Enterprise Linux 9.0	X64	Red Hat	Filesystem: ext3, ext4, xfs. Kernel version: 5.14.0-70
Red Hat Enterprise Linux 9.1	X64	Red Hat	Filesystem: ext3, ext4, xfs. Kernel version: NA
Red Hat Enterprise Linux 9.2	X64	Red Hat	Filesystem: ext3, ext4, xfs. Kernel version: NA
SLES 12 SP3	X64	SUSE	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 4.4.73-5
SLES 12 SP4	X64	SUSE	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 4.12.14-94.41
SLES 12 SP5	X64	SUSE	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 4.12.14-120
SLES 15 SP1	X64	SUSE	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 4.12.14-195
SLES 15 SP2	X64	SUSE	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 5.3.18-22
SLES 15 SP3	X64	SUSE	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 5.3.18-57
SLES 15 SP4	X64	SUSE	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 5.14.21-150400
CentOS 7.0	X64	CentOS	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 3.10.0-123
CentOS 7.1	X64	CentOS	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 3.10.0-229
CentOS 7.2	X64	CentOS	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 3.10.0-327
CentOS 7.3	X64	CentOS	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 3.10.0-514
CentOS 7.4	X64	CentOS	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 3.10.0-693
CentOS 7.5	X64	CentOS	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 3.10.0-862
CentOS 7.6	X64	CentOS	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 3.10.0-957
CentOS 7.7	X64	CentOS	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 3.10.0-1062
CentOS 7.8	X64	CentOS	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 3.10.0-1127

CentOS 7.9	X64	CentOS	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 3.10.0-1160.
CentOS 8.0	X64	CentOS	Filesystem: ext3, ext4, xfs. Kernel version: 4.18.0-80.
CentOS 8.1	X64	CentOS	Filesystem: ext3, ext4, xfs. Kernel version: 4.18.0-147.
CentOS 8.2	X64	CentOS	Filesystem: ext3, ext4, xfs. Kernel version: 4.18.0-193.
CentOS 8.3	X64	CentOS	Filesystem: ext3, ext4, xfs. Kernel version: 4.18.0-240.
Ubuntu 16.04	X64	Canonical	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 4.4.0-21.
Ubuntu 18.04	X64	Canonical	Filesystem: ext3, ext4, xfs, btrfs. Kernel version: 4.15.0-29.
Ubuntu 20.04.1 LTS	X64	Canonical	Filesystem: ext3, ext4, xfs. Kernel version: 5.8.0-59-generic.
Oracle Linux 7.x UEK	X64	Oracle	Filesystem: ext3, ext4, xfs. Kernel version: NA.
Oracle Linux 8.x UEK	X64	Oracle	Filesystem: ext3, ext4, xfs. Kernel version: NA.
Oracle Linux 9.x UEK	X64	Oracle	Filesystem: ext3, ext4, xfs. Kernel version: NA.
AIX 7.1 TL5 SP2	X64	IBM	Filesystem: JFS, JFS2. Kernel version: NA.
AIX 7.2 TL2 SP1	X64	IBM	Filesystem: JFS, JFS2. Kernel version: NA.
AIX 7.2 TL3	X64	IBM	Filesystem: JFS, JFS2. Kernel version: NA.
AIX 7.2 TL4	X64	IBM	Filesystem: JFS, JFS2. Kernel version: NA.
AIX 7.2 TL5	X64	IBM	Filesystem: JFS, JFS2. Kernel version: NA.
AIX 7.3 TL0	X64	IBM	Filesystem: JFS, JFS2. Kernel version: NA.
AIX 7.3 TL1	X64	IBM	Filesystem: JFS, JFS2. Kernel version: NA.

Notes:

- The File System agent supports Microsoft failover clustering for both CSV and cluster volumes configurations.
- Kernel versions are applicable for BBB (block-based backups), that is, for image-level backups ONLY. The File System agent uses BBB technology if there are no exclusion filters associated with a policy. Only LVM-based volumes are supported for protection. Backups of non-LVM/physical disks are not supported with BBB.
- If a protection policy is associated with one or more exclusion filters, the file system agent performs backups using FBB (file-based backup) technology. Backups using FBB are independent of the kernel version.
- Backups in the following scenarios are always taken using FBB regardless of whether the policy is associated with exclusion filters or not, as BBB does not support these use cases:
 - Backups on Oracle Linux 7.x UEK platform (use `--skip-driver` during agent installation to skip the BBB driver)
 - Backups on Ubuntu 20.04.1 LTS
 - Backups of non-LVM/physical disks
 - Backups of btrfs volumes
- The File System agent on AIX supports only FBB. BBB is not supported.
- The File System agent on AIX requires the `xlc` runtime environment 16.1.0.7 or later.
- The *PowerProtect Data Manager for File System Agent User Guide* provides installation steps and configuration information for AIX.
- The File System agent supports disaster recovery backups, which consist of bare metal recovery (BMR) backup and system state recovery (SSR) backup, for Windows only.
- The File System agent supports Microsoft Distributed File System (DFS).

POWERPROTECT DATA MANAGER IN CLOUD MARKETPLACES

Cloud Provider	Cloud Service	Supported Assets	Supported Regions
Amazon	AWS, AWS Glacier Instant Retrieval (IR)	SQL, Oracle, SAP HANA, and File System agents, and Kubernetes asset sources are supported. VM protection is available for VMs in VMC for AWS (SDOC 1.16, 1.18, 1.20).	[DDOS 7.4.0.5 onwards]: ca-central-1 (Canada - Central), eu-south-1 (EU - Milan), me-south-1 (Middle East - Bahrain) [From DDOS 6.0 onwards]: us-east-1, us-west-1, us-west-2, eu-west-1, ap-northeast-1, ap-southeast-1, ap-southeast-2, sa-east-1 (South America - Sao Paulo), ap-south-1 (Asia Pac - Mumbai), ap-northeast-2 (Asia Pac - Seoul), eu-central-1 (EU - Frankfurt)
	AWS GovCloud/C2S	SQL, Oracle, SAP HANA, and File System agents, and Kubernetes asset sources are supported. VM protection is available for VMs in VMC for AWS (SDOC 1.16, 1.18, 1.20).	[DDOS 7.2 onwards]: us-gov-east-1, us-gov-west-1
Microsoft	Azure	SQL, Oracle, SAP HANA, and File System agents, and Kubernetes asset sources are supported. VM protection is available for VMs in Azure VMWare Solution.	All regions listed under General Purpose storage class. Azure Gov requires DDOS 6.1.1.x or later.
	Azure Government Cloud	SQL, Oracle, SAP HANA, and File System agents, and Kubernetes asset sources are supported. VM protection is available for VMs in Azure VMWare Solution.	US Gov Arizona, US Gov Texas, US Gov Virginia
Google	GCP	SQL, Oracle, SAP HANA, and File System agents, and Kubernetes asset sources are supported. VM protection is available for VMs in Google Cloud VMware Engine.	us, eu, asia, northamerica-northeast1, us-central1, us-east1, us-east4, us-west1, southamerica-east1, europe-north1, europe-west1, europe-west2, europe-west3, europe-west4, asia-east1, asia-northeast1, asia-south1, asia-southeast1, australia-southeast1

Notes:

- Enabling Cloud DR and Search functionality is not supported in PowerProtect Data Manager instances running in the public cloud. This restriction does not apply to PowerProtect Data Manager instances running in VMware in Cloud.

CLOUD DISASTER RECOVERY

VMware Infrastructure

Platform	Platform Version	Comments
VMware vCenter/ESXi	6.5, 6.7, 7.0, 7.0u1, 7.0u2, 7.0u3, 8.0	

Cloud Vendor Support

Vendor Name	Cloud Service
Amazon	AWS (Amazon Web Services) & AWS GovCloud
Microsoft	Microsoft Azure & Azure Government Cloud

AWS Target Supported Virtual Machine Operating Systems

Windows 64-bit	Linux or Unix
Microsoft Windows Server 2012 Standard, Datacenter. 64-bit only.	Ubuntu 14.04, 14.10, 16.04, 16.10
Microsoft Windows Server 2012 R2 Standard, Datacenter. 64-bit only. Nano Server installation not supported.	Red Hat Enterprise Linux 6.1–6.9, 7.0–7.3, 8 (6.0 lacks required drivers)
Microsoft Windows Server 2016 Standard, Datacenter. 64-bit only.	SUSE Linux Enterprise Server 15 SP3 and SP4
Microsoft Windows Server 2019 Standard, Datacenter. 64-bit only.	CentOS 7
Microsoft Windows 10 Professional, Enterprise, Ultimate. US English. 32-and 64-bit.	Oracle Linux 6.1–6.6, 7.0–7.1

Supported AWS Regions

Region Name	Region ID	Supported
US East (Ohio)	us-east-2	Yes
US East (N. Virginia)	us-east-1	Yes
US West (N. California)	us-west-1	Yes
US West (Oregon)	us-west-2	Yes
Africa (Cape Town)	af-south-1	Yes
Asia Pacific (Hong Kong)	ap-east-1	Yes
Asia Pacific (Mumbai)	ap-south-1	Yes
Asia Pacific (Singapore)	ap-southeast-1	Yes
Asia Pacific (Sydney)	ap-southeast-2	Yes
Asia Pacific (Jakarta)	ap-southeast-3	Yes (Windows only)
Asia Pacific (Tokyo)	ap-northeast-1	Yes
Asia Pacific (Seoul)	ap-northeast-2	Yes
Asia Pacific (Osaka)	ap-northeast-3	Yes
Canada (Central)	ca-central-1	Yes
Europe (Frankfurt)	eu-central-1	Yes
Europe (Ireland)	eu-west-1	Yes
Europe (London)	eu-west-2	Yes
Europe (Paris)	eu-west-3	Yes
Europe (Stockholm)	eu-north-1	Yes
Europe (Milan)	eu-south-1	Yes
Middle East (Bahrain)	me-south-1	Yes
South America (Sao Paulo)	sa-east-1	Yes

Supported AWS GovCloud Regions

AWS GovCloud (US-Gov-West)	us-gov-west-1	Yes
AWS GovCloud (US-Gov-East)	us-gov-east-1	Yes

Azure Cloud Supported Virtual Machine Operating Systems

Windows 64-bit	Linux or Unix
Microsoft Windows Server 2012 Standard, Datacenter. 64-bit only.	Ubuntu 14 and 16
Microsoft Windows Server 2016 Standard, Datacenter. 64-bit only.	Red Hat Enterprise Linux 6-9 (DR supported only for the registered machines)
Microsoft Windows 10 Professional, Enterprise, Education. US English 64-bit	CentOS 7
Microsoft Windows Server 2019 Standard, Datacenter. 64-bit only.	Oracle Linux 6.4 - 7 (non-UEK3 and up)
Microsoft Windows Server 2022 Standard, Datacenter. 64-bit only.	Debian 9, 10, 11. SLES 15 SP3 and SP4.

Supported Azure Regions

Region Name	Region ID	Supported
Central US	us-central	Yes
East US	us-east	Yes
East US 2	us-east-2	Yes
North Central US	us-north-central	Yes
South Central US	us-south-central	Yes
West Central US	us-west-central	Yes
West US	us-west	Yes
West US 2	us-west-2	Yes
North Europe	europa-north	Yes
West Europe	europa-west	Yes
East Asia	asia-pacific-east	Yes
Southeast Asia	asia-pacific-southeast	Yes
Japan East	japan-east	Yes
Japan West	japan-west	Yes
Brazil South	brazil-south	Yes
Australia East	australia-east	Yes
Australia Southeast	australia-southeast	Yes
Central India	central-india	Yes
South India	south-india	Yes
West India	west-india	Yes
Canada Central	canada-central	Yes
Canada East	canada-east	Yes
UK South	united-kingdom-south	Yes
UK West	united-kingdom-west	Yes
France Central	france-central	Yes
France South	france-south	Yes
Korea Central	korea-central	Yes
Korea South	korea-south	Yes
South Africa North	south-africa-north	Yes
UAE North	uae-north	Yes

Supported Azure Government Cloud Regions

Region Name	Region ID	Supported
Virginia	US-Gov-Virginia	Yes

Notes:

- Cloud DR is not supported as part of PowerProtect Data Manager when deployed in AWS, Azure, or GCP.

KUBERNETES

Distributions

Supported Distributions	Versions
All Cloud Native Computing Foundation (CNCF) certified Kubernetes distributions	1.23 – 1.27
Distributions Validated and Tested by Dell	Versions
Upstream Kubernetes	1.23 – 1.27
RedHat OpenShift	4.10, 4.11 ¹ , 4.12 ¹ (using OADP 1.1.3)
VMware Tanzu Kubernetes Grid Integrated (TKGi)	1.14, 1.15, 1.16, 1.17
VMware TKGm (Multi-Cloud) Workload Cluster/TKG 2.x (Standalone)	1.5, 1.6, 2.1.1, 2.2
VMware vSphere with Tanzu Guest Cluster (TKGs/TKG 2.0)	7.0 U3(d-h), 8.0, 8.0 U1 ²
SUSE Rancher Kubernetes Engine (RKE)	1, 2
SUSE K3s on SLE Micro	
Google Anthos (On-Premises)	

Notes:

- Restore to another cluster is not supported in OpenShift for namespaces using integrated registry.
 - vSphere with Tanzu Supervisor Clusters is not supported for protection.
1. vSphere CSI driver storage classes with volume binding mode of `WaitForFirstConsumer` are not supported (for all distributions).
 2. Includes the vSphere Distributed Switch (VDS) networking stack. Requires `velero-vsphere-operator` 1.4.

Cloud Distributions

Distributions Validated and Tested by Dell	Versions
Microsoft Azure Kubernetes Service (AKS)	
Google Kubernetes Engine (GKE)	
Amazon Elastic Kubernetes Service (EKS)	

Container Images

Image	Repository	Tag
dell EMC/powerprotect-k8s-controller	https://hub.docker.com/r/dell EMC/powerprotect-k8s-controller/tags	19.14.0-20
dell EMC/powerprotect-cproxy	https://hub.docker.com/r/dell EMC/powerprotect-cproxy/tags	19.14.0-20
dell EMC/powerprotect-velero-dd	https://hub.docker.com/r/dell EMC/powerprotect-velero-dd/tags	19.14.0-20
velero/velero	https://hub.docker.com/r/velero/velero/tags	v1.9.3
vmware-veleroplugin/velero-plugin-for-vsphere	https://hub.docker.com/r/vsphereveleroplugin/velero-plugin-for-vsphere/tags	v1.5.0
vsphereveleroplugin/backup-driver	https://hub.docker.com/r/vsphereveleroplugin/backup-driver/tags	v1.5.0

Storage

Supported Storage

Any on-prem or cloud storage Container Storage Interface (CSI) driver with snapshot support, version 1.0 or higher.

- List of available CSI drivers: <https://kubernetes-csi.github.io/docs/drivers.html>
- CSI support matrix for Dell Technologies storage: <https://dell.github.io/csm-docs/docs/csdriver/>

CSI Drivers Validated and Tested by Dell	vSphere Versions	Notes
vSphere Cloud Native Storage (CNS) storage for Tanzu Kubernetes guest clusters	7.0u3(d-h), 8.0	vSphere CSI driver 2.5 to 3.0 using FCD snapshots ^{3,4}
vSphere CNS for native K8s clusters	6.7u3 and above	vSphere CSI driver 2.5 to 3.0 using FCD snapshots ^{3,4}
OpenShift Data Foundation (ODF) <ul style="list-style-type: none"> CEPH-FS¹ CEPH-RDB 		
PowerFlex		
PowerScale ²		
PowerStore		

Notes:

- For NFS-based CSI storage, storage class with root client access enabled is required for backup. See the *PowerProtect Data Manager for Kubernetes User Guide* for information about protection of root squashed persistent volumes.
- Limited support for CEPH-FS persistent volumes. Protection may fail for volumes greater than 25 GB.
 - CSI driver version for PowerScale must be 1.5 or higher and the user provided must additionally have `ISI_PRIV_IFS_BACKUP` and `ISI_PRIV_IFS_RESTORE` privileges. See *CSI Driver for PowerScale Installation Guide* for more information. The storage classes in the Kubernetes cluster should use only one PowerScale access zone.
 - vSphere CSI driver storage classes with volume binding mode of `WaitForFirstConsumer` are not supported.
 - vSAN File Services are not supported.

SUPPORTASSIST

Gateway	Gateway Version
SCG	5.10 and higher
Notes: <ul style="list-style-type: none"> SRS and SAE gateways are not supported by SupportAssist 4.0 and higher. Update these gateways to SCG before you update PowerProtect Data Manager. Alternatively, you can configure SupportAssist by connecting directly to the backend. 	

SYSLOG FORWARDING

Protocol	Rsyslog Version
UDP	8.24.0 - 8.2204.0
TCP	8.24.0 - 8.2204.0
TLS	8.24.0 - 8.40.0
Notes: <ul style="list-style-type: none"> When you use TLS, the installed versions of <code>rsyslog-gnutls</code> and <code>rsyslog</code> must match. 	