

Dell PowerVault ME5 Series

Administrator's 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.

Chapter 1: Getting started	8
ME5 Series system requirements.....	8
Web browser requirements.....	8
SupportAssist direct connection requirements.....	9
Product features.....	9
Configuring and provisioning a new storage system.....	10
Using PowerVault Manager	11
Areas of the interface.....	11
Tips for using the PowerVault Manager.....	13
Tips for using tables.....	13
Exporting data to a CSV file.....	13
Size representations.....	13
Signing in and signing out.....	15
Chapter 2: System concepts	16
Storage environment.....	17
Virtual storage.....	17
Linear storage.....	17
Disk groups.....	17
Disk sector format.....	18
Virtual disk group configuration.....	18
Removing disk groups from virtual pools.....	19
Linear disk group configuration.....	19
Read-cache disk groups.....	20
RAID levels.....	20
ADAPT.....	22
Disk group utilities.....	23
Disk group expansion.....	23
Disk-group scrub.....	23
SSDs.....	24
Gauging the percentage of life remaining for SSDs.....	24
Viewing I/O workload activity.....	24
SSD management.....	25
SSD read cache.....	26
Spares.....	27
Pools.....	27
Virtual pools and disk groups.....	27
Linear pools and disk groups.....	28
Changing pool settings.....	28
Volumes and volume groups.....	29
Virtual volumes.....	29
Volume groups.....	29
Linear volumes.....	29
Volume cache options.....	29

Using write-back or write-through caching.....	30
Cache optimization mode.....	30
Optimizing read-ahead caching.....	30
Overcommitting volumes.....	31
About automated tiered storage.....	31
Volume tier affinity.....	31
Initiators, hosts, and host groups.....	32
CHAP.....	32
Host ports.....	33
Attaching volumes to hosts.....	33
Operating with a single controller.....	33
Snapshots.....	34
Snapshot trees.....	34
Rollback and reset snapshot features.....	34
Copying volumes or snapshots.....	35
Reconstruction.....	35
ADAPT reconstruction.....	35
Quick rebuild.....	36
Updating firmware.....	36
Managed logs.....	36
Saving log data to a file.....	37
DNS settings.....	38
Peer connections.....	38
Replication.....	39
Replication prerequisites.....	39
Replication process.....	40
Creating a virtual pool for replication.....	42
Setting up snapshot space management in the context of replication.....	42
Replication and empty allocated pages.....	43
Disaster recovery.....	43
Creating a replication set.....	45
CHAP and replication.....	46
Full disk encryption	47
Rescanning disks.....	48
Clearing disk metadata.....	48
Data protection during failover to a single controller.....	49
SupportAssist.....	49
SupportAssist data.....	49
About CloudIQ.....	50
Event history.....	50
Audit logs.....	50
System metrics.....	51
Effect of failover on metrics.....	51
Chapter 3: Dashboard.....	52
Alerts panel.....	52
Compact view.....	52
Expanded view.....	52
Capacity panel.....	53
Compact view.....	53

Expanded view.....	54
Performance panel.....	54
Compact view.....	54
Expanded view.....	55
Performance metrics.....	55
Activity panel.....	57
Compact view.....	57
Expanded view.....	57
Chapter 4: Provisioning.....	58
Working with volumes.....	58
Volumes table.....	58
Data protection table.....	59
Creating volumes.....	60
Modifying volumes.....	61
Deleting volumes and snapshots.....	61
Attaching volumes to hosts.....	61
Detaching volumes from hosts.....	61
Expanding volumes.....	62
Rolling back virtual volumes.....	62
Creating snapshots.....	62
Resetting snapshots.....	62
Copying volumes or snapshots.....	63
Aborting a volume copy.....	63
Add data protection.....	63
Creating a replication set.....	63
Modifying a replication set.....	64
Deleting a replication set.....	64
Initiating or scheduling a replication.....	65
Suspending a replication.....	65
Aborting a replication set.....	66
Resuming a replication.....	66
Managing replication schedules.....	66
Working with hosts.....	66
Creating hosts.....	67
Attaching hosts to volumes.....	67
Detaching hosts from volumes.....	67
Removing initiators from a host.....	67
Removing hosts from a host group.....	68
Adding hosts to a host group.....	68
Deleting hosts.....	68
Deleting host groups.....	68
Add initiators to a host.....	68
Renaming hosts.....	69
Changing a host profile.....	69
Renaming a host group.....	69
Renaming initiators.....	69
Chapter 5: Settings.....	70

Network settings.....	70
Configuring controller network ports.....	70
Configuring DNS settings.....	71
Enabling or disabling system-management services.....	71
Viewing certificate information.....	72
User settings.....	72
Managing local users.....	72
Managing SNMPv3 users.....	73
System settings.....	74
Setting system identification information.....	74
Setting the date and time.....	74
Securing the system with FDE.....	74
Setting system properties.....	75
Notification settings.....	76
Email notifications.....	76
SNMP notifications.....	77
Syslog notifications.....	77
iSCSI host port settings.....	77
Configuring iSCSI CHAP settings.....	77
Changing iSCSI configuration settings.....	78
Peer connection settings.....	78
Querying peer connections.....	78
Modifying peer connection settings.....	78
Deleting a peer connection.....	79
Chapter 6: Maintenance.....	80
Storage panel.....	80
Viewing information about a pool.....	81
Changing pool settings.....	81
Viewing information about volumes for each disk group in a pool.....	81
Adding a disk group to a pool.....	82
Renaming a disk group.....	82
Deleting a disk group from a pool.....	82
Expanding an ADAPT disk group.....	82
Scrubbing a disk group.....	82
Managing spares.....	82
Hardware panel.....	82
Firmware panel.....	84
Viewing information about installed and active system firmware bundles.....	84
Updating system firmware.....	84
Updating disk firmware.....	85
Best practices for updating firmware.....	86
About panel.....	86
Support panel.....	87
Configuring SupportAssist.....	87
Enable SupportAssist.....	88
Enable or disable CloudIQ.....	90
Appendix A: Other management interfaces.....	91


SNMP reference.....	91
Supported SNMP versions.....	91
Standard MIB-II behavior.....	91
Enterprise traps.....	92
FA MIB 2.2 SNMP behavior.....	92
External details for certain FA MIB 2.2 objects.....	97
Configuring SNMP event notification in the PowerVault Manager.....	99
SNMP management.....	100
Using FTP and SFTP.....	100
Downloading system logs.....	100
Transferring log data to a log-collection system.....	101
Downloading historical disk-performance statistics.....	102
Updating firmware.....	103
Installing a security certificate.....	107
Download a system heat map data.....	108
Using SLP.....	109
Appendix B: Administering a log-collection system.....	111
How log files are transferred and identified.....	111
Log file details.....	111
Storing log files.....	112
Appendix C: Settings changed by restoring defaults.....	113
Appendix D: System configuration limits.....	117
Appendix E: Best practices.....	120
Pool setup.....	120
RAID selection.....	120
Disk count per RAID level.....	121
Disk groups in a pool.....	121
Tier setup.....	122
Multipath configuration.....	122
Appendix F: Glossary of terms.....	124

Getting started

The PowerVault Manager is a web-based application for configuring, monitoring, and managing the storage system. The PowerVault Manager is also referred to as the web-browser interface (WBI).

Each controller module in the storage system contains a web server, which is accessed when you sign in to the PowerVault Manager. You can access all functions from either controller in a dual-controller system. If one controller becomes unavailable, you can continue to manage the storage system from the partner controller.

In addition to the PowerVault Manager, each controller module in the storage system has the following interfaces: SNMP, FTP, SFTP, SLP, CLI, API. For information about using the CLI and API, see the CLI Reference Guide.

 **NOTE:** SMI-S is not supported on the ME5 Series storage systems.

Topics:

- [ME5 Series system requirements](#)
- [Product features](#)
- [Configuring and provisioning a new storage system](#)
- [Using PowerVault Manager](#)

ME5 Series system requirements

The following sections list browser and network requirements for the ME5 Series system.


Web browser requirements

PowerVault Manager supports the browsers that are listed below.

- Apple Safari 11 and newer (Mac)
- Google Chrome 70 and newer
- Microsoft Internet Explorer 11
- Mozilla Firefox 68 and newer

For best results, follow these guidelines:

- The recommended resolution for the page display area in the browser is 1360 x 768 pixels.
- To optimize the display, use a color monitor and set its color quality to the highest setting.
- To navigate beyond the sign-in page (with a valid user account):
 - If the PowerVault Manager is configured to use HTTPS, ensure that your browser is set to use TLS 1.2.
 - Verify that the browser is set to allow cookies, at least for the IP addresses of the storage system network ports.
 - For Internet Explorer, set the local-intranet security option in the browser to medium or medium-low.
 - For Internet Explorer, add the network IP address of each controller module as a trusted site.
- To see the help window, enable pop-up windows.

 **NOTE:** By default, your system is loaded with self-signed certificates. Generate new self-signed certificates on each controller using the `create_certificate` CLI command. Browser messages warning you about security or privacy concerns due to self-signed or untrusted certificates or invalid certificate authorities are expected, and warnings can be safely bypassed if you trust that you are contacting the correct controller within your network. Depending on the browser and its settings, once you navigate through the browser warning, a security exception may be created and, the warning may not appear. Your browser address bar still indicates that the connection is not trusted or not secure. You can safely ignore the indication if you trust you are accessing the correct controller within your network.

SupportAssist direct connection requirements

The following network requirements must be met to use a direct connection for SupportAssist.

- A minimum of one DNS server must be configured.
- The local system must connect to the following destinations to ensure connectivity to the Global servers:
 - esrs3-core.emc.com
 - esrs3-core.dr.emc.com

Use the following command to verify the connections:

```
# check support-assist-connection mode direct
```

If the connection is working, it returns an HTTP Status 200 message:

```
SupportAssist Connection
-----
Connection State: Connected

Endpoint
-----
Mode: direct
Endpoint: https://esrs3-core.emc.com
Status: success
HTTP Status: 200
Message: OK
Status Detail: Success
Proxy Type: none

Endpoint
-----
Mode: direct
Endpoint: https://esrs3-core.dr.emc.com
Status: success
HTTP Status: 200
Message: OK
Status Detail: Success
Proxy Type: none
Success: Command completed successfully. (2022-01-08 18:04:00)
```

To ensure connection integrity, proxy servers and devices external to your demilitarized zone (DMZ) must not perform any method of SSL decryption on outbound or inbound traffic for the Dell secure remote servers. SSL decryption that is performed on outbound communication causes a loss of connectivity to the backend. SSL decryption includes decryption by your firewall, proxies, web traffic filtering appliances or cloud services, web traffic shaping or load balancing, certificate verification, certificate proxy, or Intrusion Detection Services (IDS).

In case the SSL decryption is enabled on the proxy servers and other devices, ensure that `esrs3-core.emc.com` and `esrs3-core.dr.emc.com` are added to the SSL decryption exclusion list on the proxy servers and devices.

Table 1. Port requirements

Type of connection	Ports that must be open	Protocol used	Communication
Connect Directly	443	TCP	Outbound
Connect via Gateway Server	9443	TCP	Outbound

Product features

PowerVault Manager provides access to many features that help you manage the storage system.

Some of the key features include:

- **Storage system setup**—the capability to initially connect to a system using PowerVault Manager. PowerVault Manager uses intuitive preboarding and onboarding steps to guide you through initial setup of your storage, as described in [Configuring and provisioning new storage](#) system.

- **SupportAssist**—the ability to monitor and proactively maintain the health of your ME5 Series storage system. SupportAssist is also required to use CloudIQ, which is a no-cost cloud-based storage analytics application.
- **ADAPT data protection**—RAID-based data protection level that emphasizes efficiency, as described in [ADAPT](#).
- **Replication**—the capability to replicate block-level data from a volume in a primary system to a volume in a secondary system, as described in [Replication](#).
- **Update firmware**—the capability to notify users of available firmware updates to controller modules, expansion modules, and disk modules with newer and compatible firmware versions as they become available, as described in [Updating firmware](#).
- **Performance metrics**—the ability to monitor storage system performance and statistics using graphs generated live or from historical data as described in [Performance panel](#).
- **Alerts**—a robust storage enclosure health and notification system that is designed to identify actionable conditions and promote best practices, as described in [Alerts panel](#).
- **SSDs**—the ability to use solid state drives to enhance storage system performance, as described in [SSDs](#). The ability to move data from one class of disks to another automatically, based on data access patterns, as described in [Automated tiered storage](#).
- **Virtual storage**—a storage model that maps logical components to physical media-using paged-storage technology-to virtualize data storage, as described in [Virtual storage](#).
- **Linear storage**—a storage model that maps logical components to physical media, as described in [Linear storage](#).
- **IPv6 support**—the capability for the storage system to support Internet Protocol version 6 (IPv6) functionality (in addition to IPv4), as described in [Configuring controller network ports](#).
- **Redfish REST API support**—the Redfish REST (Representational State Transfer) API provides the management data in a stateless, cacheable data representation. Read-only access is provided to physical and logical components that are related to the storage provisioning model, including disks, storage pools, volumes, and enclosures. The public APIs called DMTF Redfish and SNIA Swordfish are supported:
 - For technical information about DMTF Redfish, see <https://www.dmtf.org/standards/redfish>.
 - For technical information about SNA Swordfish, see <https://www.snia.org/forums/smi/swordfish>.
 - The base URL for accessing the Redfish API capability is "https://<controller-module-IP-address>/redfish". (Enter the unquoted portion of the URL string into your browser address field using a valid controller-module-IP-address in place of the variable text.)
 - To obtain an open source cross-platform Python tool for provisioning and managing storage systems using the RESTful Redfish/Swordfish API see <https://github.com/Seagate/SystemsRedfishPy>.
- VMware SRA and SRM support—enables the ability to integrate the Storage Replication Adapter (SRA) with the Site Recovery Manager (SRM) to manage replication features for the storage system using the management interfaces.

Configuring and provisioning a new storage system

When you connect to the system for the first time, a wizard in the PowerVault Manager guides you through the first-time setup of your system. This process is called preboarding and onboarding.

See the *ME5 Series Storage System Deployment Guide* for details about the options that are presented in the guided setup wizard.

During preboarding you are led through steps to prepare the system for use and are prompted to do the following:

- Agree to the EULA.
- Create a username, password, and select a language (once complete, you are logged into the system as this user).
 - **NOTE:** The user that is created during the preboarding process has managing capabilities and can change system settings.
- Update firmware

During onboarding, you are led through steps to configure and provision the system. These steps include:

- Configuring system settings:
 - Network settings (IPv4, IPv6, DNS)
 - Date and time (NTP)
 - User definitions (local, SNMPv3)
 - Notifications (email, SNMP, syslog)
- Configuring SupportAssist:
 - License
 - Connection
 - Support contacts
 - Summary

- Configuring storage settings:
 - Storage type (Linear|Virtual)
 - Automatic disk group and pool creation
- Manual (advanced) disk group and pool creation
- Provisioning storage:
 - Creating hosts and host groups (naming initiators, assigning initiators to hosts, creating a single host)
 - Creating volumes and attaching them to hosts
 - Summary
 - Repeat

Follow the on-screen directions to complete setting up your system. Once you complete the preboarding and onboarding steps the PowerVault Manager Dashboard opens where you monitor, manage, and provision the storage system.

Using PowerVault Manager

This section describes the PowerVault Manager interface and provides tips for using it.

Areas of the interface

The main areas of the PowerVault Manager interface are the banner, the menu pane, and the management pane, as represented by the following table. For more information about an item in the banner or menu pane, click its name in the table.

- **Banner**—The banner provides links to the PowerVault Manager Help, [Setting the date and time](#), [User settings](#), and the Log Out option.
- **Menu pane**—Clicking an option on the menu pane on the left side expands a dropdown list of menu choices. Clicking a menu option displays applicable content in the management pane. The menu pane options include [Dashboard](#), [Provisioning](#), [Settings](#), and [Maintenance](#).
- **Management pane**—The management pane shows system status relating to the selected menu in a summary format, allowing you to monitor and interact with the system. Where applicable, you can expand summary sections by clicking the slide-over arrows to view more information about the system status and make applicable changes to system settings and configuration. You can click on the information icon to view content that defines or explains more information about a feature/option.

Icons in the interface

The following table displays a list of the most common icons found in the PowerVault Manager.

Table 2. PowerVault Manager icons











Icon	Name	Use
	Abort/Cancel	Aborts or cancels an operation.
	Apply	Applies an edited operation or selection.
	Critical	Indicates that the item's health is critical, or that an alert or event has Critical severity.
	Collapse	Collapses a table row to hide information about an object.
	Expand	Expands a table row to provide more detail about an object.
	Dashboard	Indicates that the Dashboard menu is selected.
	Degraded	Indicates that the item's health is degraded, or that an alert or event has Error severity.
	Delete	Lets you delete a value or object.
	Disk	Indicates an operation was performed on a disk.

Table 2. PowerVault Manager icons (continued)





Icon	Name	Use
	Disk group	Indicates an operation was performed on a disk group.
	Edit	Lets you edit a single value or options within an entire row or table.
	Error/Fault	Indicates that there is an error or fault with the system.
	Export/Upload	Lets you export or upload a file.
	Favorite	Indicates that the graph selected is a favorite and will display in the compact view of the dashboard.
	Healthy/OK alert	Indicates that the item's health is good, or an alert or event is resolved or acknowledged.
	Host	Identifies a host.
	Host group	Identifies a host group.
	Information	Opens a small window that defines or provides more information about a feature or option.
	Informational	Indicates that an alert or event is informational.
	Initiator	Identifies an initiator.
	Maintenance	Indicates that the Maintenance menu is selected.
	Maintenance tasks	Indicates that either maintenance needs to be performed or has already been performed to the specified item.
	Primary replication volume	Identifies the primary replication volume.
	Provisioning	Indicates that the Provisioning menu is selected.
	Resume	Resumes a suspended operation.
	Schedule	Indicates that a specified task will take place at defined times.
	Secondary replication volume	Identifies the secondary replication volume.
	Secured	Indicates that the system is secured using FDE.
	Settings	Indicates that the Settings menu is selected.
	Slide-over arrows	Opens or closes a panel that contains detailed information about an object.
	Snapshot	Indicates that a snapshot of the volume was created.
	Suspend	Suspends (pauses) an in-progress operation.
	Unsecured	Indicates that the system is not secured using FDE.
	Volume	Identifies the primary volume.
	Volume copy	Indicates the volume is being copied to a new volume.

Table 2. PowerVault Manager icons (continued)

Icon	Name	Use
	Warning	Indicates that an alert or event has Warning severity.




Tips for using the PowerVault Manager

The following list contains tips for using PowerVault Manager.

- Do not use the Back or Forward buttons in the browser. The PowerVault Manager has a single page for which content changes as you perform tasks, and automatically updates to show current data.
- If you are signed in to the PowerVault Manager and the controller you are accessing goes offline, the system informs you that the system is unavailable or that communication has been lost. After the controller comes back online, close and reopen the browser and start a new PowerVault Manager session.
- As you set options in panels, the PowerVault Manager informs you whether a value is invalid or a required option is not set.
- Confirmation buttons become active only after you set all required options.
- A red asterisk (*) identifies a required setting.
- Click  to expand a panel and view additional details or perform actions. Click  to collapse a panel and view summary information.
- Click  to open the information window to learn more about an option. Click the icon again to close the information window.
- Click  to edit content within a text box or table.
- In the Hardware panel (**Maintenance > Hardware**), click on a component such as an enclosure or disk to display information about that component.

Tips for using tables

Use the following controls to work with PowerVault Manager tables.

- Click  to expand a table and see additional details or perform actions. Click  to collapse a table and hide detailed information.
- The presence of a slide-over arrow icon  at the end of a table row indicates that you can view more information about the option and perform actions.
- Use the Search bar in the table header to search for specific content within the table. Not all tables have a search option.
- Table items are sorted by the highlighted column heading.
- To sort items by a specific column, click the arrow icon in the column heading to reorder items from low to high. Click the arrow icon again to reorder items from high to low.
- To filter items in a table, select the filter content from the Filter By drop-down list. Not all tables have a filtering option.
- To select items in a table, use the check boxes in the left column. Clear the check boxes to deselect items.

Exporting data to a CSV file

You can export performance data to a downloadable comma-separated values (CSV) file that you can view in a spreadsheet for further analysis. The exported CSV file contains all of the content in the graph.



To export performance data, click . For more information, see:

- [Downloading historical disk performance statistics](#)
- [Downloading system heat map data](#)

Size representations

Parameters such as names of users and volumes have a maximum length in bytes.

When encoded in UTF-8, a single character can occupy multiple bytes. Typically:

- 1 byte per character for English, Dutch, French, German, Italian, and Spanish
- 3 bytes per character for Chinese, Japanese, and Korean

Operating systems usually show volume size in base-2. Disks usually show size in base-10. Memory (RAM and ROM) size is always shown in base-2. In the PowerVault Manager, the base for entry and display of storage-space sizes can be set per user. When entering storage-space sizes only, either base-2 or base-10 units can be specified. The locale setting determines the character used for the decimal (radix) point.

Table 3. Storage size representations in base 2 and base 10

Base-2		Base-10	
Unit	Size in bytes	Unit	Size in bytes
KiB (kibibyte)	1,024	KB (kilobyte)	1,000
MiB (mebibyte)	1,024 ²	MB (megabyte)	1,000 ²
GiB (gibibyte)	1,024 ³	GB (gigabyte)	1,000 ³
TiB (tebibyte)	1,024 ⁴	TB (terabyte)	1,000 ⁴
PiB (pebibyte)	1,024 ⁵	PB (petabyte)	1,000 ⁵
EiB (exbibyte)	1,024 ⁶	EB (exabyte)	1,000 ⁶

The locale setting determines the character used for the decimal (radix) point, as shown in the following table:

Table 4. Decimal (radix) point character by locale


Language	Character	Examples
English, Chinese, Japanese, Korean	Period (.)	146.81 GB 3.0 Gb/s
Dutch, French, German, Italian, Spanish	Comma (,)	146,81 GB 3,0 Gb/s

Signing in and signing out

Multiple users can be signed in to each controller simultaneously.

About this task

For each active PowerVault Manager session, an identifier is stored in the browser. Depending on how your browser treats this session identifier, you might be able to run multiple independent sessions simultaneously. For example, each instance of Internet Explorer can run a separate PowerVault Manager session. However, all instances of Firefox, Chrome, and Safari share the same session.

 **NOTE:** If the initial user has not been created, see [Configuring and provisioning a new storage system](#) for directions on signing in to the system for the first time. Otherwise, see the following procedure.

Steps

1. In the web browser address field, type `https://<IP address of a controller network port>` and press **Enter**. (Do not include a leading zero in an IP address. For example, enter 10.1.4.33 and not 10.1.4.033.)

The Sign In page opens. If the Sign In page is not displayed, verify that you have typed the correct IP address.

 **NOTE:** HTTPS is enabled by default. To enable HTTP, see [Enable or disable system management settings](#).

2. Enter the username and password of an authorized user.
 - **Username**—A username is case sensitive and can have a maximum of 29 bytes. The name cannot already exist in the system, include spaces, or include any of the following: " , < \
 - **Password**—A password is case sensitive and can have from 8 to 32 printable characters. If the password contains only ASCII characters, it must contain at least one uppercase character, one lowercase character, one numeric character, and one non-alphanumeric character. A password can include UTF-8 characters except for the following: a space or " ' , < > \
3. To display the interface in a language different than the one configured for the user, select the language from the language drop-down menu.

Language preferences can be configured for the system and for individual users. The default language is English.
4. Click **Log In**.

The System Configuration Main page or the PowerVault Manager Dashboard is displayed.

When you are ready to end your session, click **Log Out** in the banner. Do not simply close the browser window.

System concepts

This section provides overviews of system features and concepts.

Topics:

- [Storage environment](#)
- [Virtual storage](#)
- [Linear storage](#)
- [Disk groups](#)
- [RAID levels](#)
- [ADAPT](#)
- [Disk group utilities](#)
- [SSDs](#)
- [Spares](#)
- [Pools](#)
- [Changing pool settings](#)
- [Volumes and volume groups](#)
- [Volume cache options](#)
- [Overcommitting volumes](#)
- [About automated tiered storage](#)
- [Initiators, hosts, and host groups](#)
- [CHAP](#)
- [Host ports](#)
- [Attaching volumes to hosts](#)
- [Operating with a single controller](#)
- [Snapshots](#)
- [Copying volumes or snapshots](#)
- [Reconstruction](#)
- [Quick rebuild](#)
- [Updating firmware](#)
- [Managed logs](#)
- [DNS settings](#)
- [Peer connections](#)
- [Replication](#)
- [Full disk encryption](#)
- [Rescanning disks](#)
- [Clearing disk metadata](#)
- [Data protection during failover to a single controller](#)
- [SupportAssist](#)
- [About CloudIQ](#)
- [Event history](#)
- [Audit logs](#)
- [System metrics](#)

Storage environment

This product uses two different storage technologies that share a common user interface: virtual storage and linear storage.

Virtual storage

Virtual storage is a method of mapping logical storage requests to physical storage (disks). It inserts a layer of virtualization such that logical host I/O requests are mapped onto pages of storage. Each page is then mapped onto physical storage. Within each page the mapping is linear, but there is no direct relationship between adjacent logical pages and their physical storage.

A page is a range of contiguous Logical Block Addresses (LBAs) in a disk group, which is one of up to 16 RAID sets that are grouped into a pool. Thus, a virtual volume as seen by a host represents a portion of storage in a pool. Multiple virtual volumes can be created in a pool, sharing its resources. This allows for a high level of flexibility, and the most efficient use of available physical resources.

Some advantages of using virtual storage are:

- It allows performance to scale as the number of disks in the pool increases.
- It virtualizes physical storage, allowing volumes to share available resources in a highly efficient way.
- It allows a volume to be comprised of more than 16 disks.
- It enables you to easily add storage on the fly.

Virtual storage provides the foundation for data-management features such as:

- [Automated tiered storage](#)
- [Read-cache disk groups](#)
- [Replication](#)
- [Quick rebuild](#)
- [Snapshots](#)
- [Overcommitting volumes](#)
- [Copying volumes or snapshots](#)

Linear storage

Linear storage is a method of mapping logical storage requests directly to physical storage.

In some cases the mapping is one-to-one, however in most cases the mapping is across groups of physical storage devices, or slices of them. This linear method of mapping is highly efficient. The negative side of linear mapping is lack of flexibility. This makes it difficult to alter the physical layout after it is established

Disk groups

A disk group is an aggregation of disks of the same type, using a specific RAID level for the purpose of storing volume data. Disk groups are used in both virtual and linear storage environments and are added to pools to configure storage.

A pool is an aggregation of one or more disk groups that serves as a container for volumes. Disk groups and pools are mutually inclusive. A pool cannot exist without at least one disk group in it, and you cannot create a disk group without selecting the pool it will reside in.


i **NOTE:** Clarification of pool and disk group terms:

- For linear storage, a storage pool can contain only one disk group; therefore the disk group is the pool, and the two terms are interchangeable.
- For virtual storage, a pool can contain multiple disk groups; therefore the pool and disk group terms are different and distinct.
- For hybrid storage, support is provided for either linear and virtual storage paradigms; so each of the above definitions applies to its corresponding storage type within the system.

All disks in a disk group must be the same type (SSD, enterprise SAS, or midline SAS). A disk group can contain disks with different capacities, sector formats, and models. If you mix disks with different capacities, the smallest disk determines the

logical capacity of all other disks in the disk group regardless of RAID levels except ADAPT. For example, the capacity of a disk group composed of one 1.2 TB disk and one 2.4 TB disk is equivalent to a disk group composed of two 1.2 TB disks. To maximize disk usage, use disks of similar size.

The PowerVault Manager configuration options that display when you are creating disk groups and pools are dependent on whether you are operating in a virtual storage environment or a linear storage environment.

 **NOTE:** After you create a disk group using one storage environment, the system will use that storage environment for additional disk groups. To switch storage environments, you must first remove all disk groups.

Related information

[Removing disk groups from virtual pools](#)

[ADAPT](#)

[Disk sector format](#)

[Virtual disk group configuration](#)

[Read-cache disk groups](#)

[Linear disk group configuration](#)

[Disk group expansion](#)

[Disk group scrub](#)


[Changing pool settings](#)

Disk sector format

The system supports 512-byte native sector size disks, 512-byte emulated sector size disks, or a mix of these sector formats.

The system identifies the sector format used by a disk, disk group, or pool as follows:

- **512n**—All disks use the 512-byte native sector size. Each logical block and physical block is 512 bytes.
- **512e**—All disks use 512-byte emulated sector size. Each logical block is 512 bytes and each physical block is 4096 bytes. Eight logical blocks will be stored sequentially in each physical block. Logical blocks may or may not be aligned with physical block boundaries.
- **Mixed**—The disk group contains a mix of 512n and 512e disks. For consistent and predictable performance, do not mix disks of different sector size types (512n, 512e).

 **CAUTION:** The emulation for 512e disks supports backward-compatibility for many applications and legacy operating systems that do not support 4K native disks. However, older versions of application software, such as virtualization software that resides between the operating system and your storage firmware, may not fully support 512e disk emulation. If not, performance degradation might result. Ensure that you have upgraded to the most recent version of any software that might be affected, and see its documentation for further information.

Virtual disk group configuration

A virtual disk group requires the selection of a pool (A or B) along with a specified number of available disks, RAID level, and spare size. If the virtual pool does not exist at the time of adding the disk group, the system will automatically create it.

Multiple disk groups (up to 16) can be added to a single virtual pool. To add disk groups to a pool (**Maintenance > Storage**), expand the Disk Groups section for the specified pool (A or B) and then select **Add Disk Group** to access the panel. Panel content is dynamic, displaying options based on protection (RAID) level selected and the available disks.

Disk group configuration requires you to select a protection level. Depending on the level selected, additional configuration options may be displayed. Only fault tolerant protection levels are available for creating virtual disk groups. Supported protection levels for virtual disk groups are: RAID 1, RAID 5, RAID 6, RAID 10, and ADAPT. If RAID 10 is specified, the disk group must have at least two subgroups.

Available disks are listed in the middle panel, and the summary panel will update as you select disks. The disk group will be added to the pool once you complete your selections and choose **Add Disk Group**.

Creating disk groups that contain only SSDs in combination with disk groups that contain only spinning disks will allow tiering within a pool. For more information, see [Automated tiered storage](#).

If a virtual disk group comprised of spinning disks is created, you can create a read-cache disk group comprised of SSDs. A pool may contain either a read-cache disk group or a virtual disk group containing SSDs, but not both. For more information, see [Read-cache disk groups](#).

NOTE: For optimal performance all virtual disk groups in the same tier should have the same RAID level, capacity disks, and physical number of disks.

If the owning controller fails, the partner controller assumes temporary ownership of the pool and resources owned by the failed controller. If the host fault-tolerant cabling configuration, with appropriate mapping, is used to connect the controllers to hosts, LUNs for both controllers are accessible through the partner controller, so I/O to volumes can continue without interruption.

Removing disk groups from virtual pools

You can remove one or more disk groups, but not all, from a virtual pool without losing data as long as there is enough space available in the remaining disk groups to move the data into.

When a virtual disk group that contains active volume data is removed, that volume data will drain (be moved) to other disk group members within the pool (if they exist). If the pool has enough space to contain the data, the disk group is removed. If the pool does not have enough space to contain the data, the disk group is not removed. When the last disk group is removed, the pool ceases to exist and will be deleted from the system automatically. Alternatively, the entire pool can be deleted, which automatically deletes all volumes and disk groups residing on that pool.

NOTE: Disk group removal (draining) can take a very long time depending on a number of factors in the system, including but not limited to: the amount of I/O traffic to the system (e.g., active I/O pages to the draining disk group); the type of the disk group page migration (enterprise SAS, midline SAS, SSD); the size of the draining disk group(s) in the system; and the number of disk groups draining at the same time.

NOTE: If you remove the last disk group in a virtual pool, a warning displays prompting you to confirm removing the disk group. If the reply is yes, the pool will be removed. If the reply is no, the disk group and the pool will remain.

Removing a disk group from a pool

In the **Maintenance > Storage** panel, locate the disk group to remove, click the icon, and follow on-screen directions.

Linear disk group configuration

A linear disk group requires the specification of a set of disks, RAID level, disk group type, and a name.

Whenever the system creates a linear disk group, it also creates an identically named linear pool at the same time. No further disk groups can be added to a linear pool.

During onboarding, you were allowed to use the Linear version of the Add Disk Group modal to configure disk groups. If you deferred the disk group configuration option during onboarding, manual configuration is available.

The Add Disk Group action (**Maintenance > Storage > Pool Configuration**) enables manual configuration of disk groups. Disk group configuration requires you to enter a specified name, assigned controller, and protection (RAID) level.

The **Add Disk Group** panel is dynamic, displaying configuration options based on the RAID level selected and the available disks on the system. Available disks are listed in the middle panel, and the summary panel will update as you select disks. The disk group will be added to the pool once you complete your selections and choose **Add Disk Group**.

The RAID levels for linear disk groups created through the PowerVault Manager must be fault tolerant. The supported RAID levels for linear disk groups in the interface are: RAID 1, RAID 5, RAID 6, RAID 10, ADAPT. RAID 10 appears in the interface only if the system disk configuration supports it. If RAID 10 is specified, the disk group has a minimum of two subgroups. Additionally, you can create non-fault-tolerant NRAID or RAID-0 disk groups through the CLI

For maximum performance, all of the disks in a linear disk group must share the same classification, which is determined by disk type, size, and speed. This provides consistent performance for the data being accessed on that disk group.

Each time that the system adds a linear disk group, it also creates a corresponding pool for the disk group. Once a linear disk group and pool exists, volumes can be added to the pool. The volumes within a linear pool are allocated in a linear/sequential way, such that the disk blocks are sequentially stored on the disk group.

Linear storage maps logical host requests directly to physical storage. In some cases the mapping is 1-to-1, while in most cases the mapping is across groups of physical storage devices, or slices of them.

To remove a linear disk group, delete the disk group and the contained volumes are automatically deleted. The disks that compose that linear disk group are then available to be used for other purposes.

If the owning controller fails, the partner controller assumes temporary ownership of the disk groups and resources owned by the failed controller. If a fault-tolerant cabling configuration, with appropriate mapping, is used to connect the controllers to hosts, LUNs for both controllers are accessible through the partner controller so I/O to volumes can continue without interruption.

Read-cache disk groups

A read-cache disk group is a special type of a virtual disk group that is used to cache virtual pages to improve read performance.

Only a single read-cache disk group may exist within a pool. Read cache does not add to the overall capacity of the pool to which it has been added. You can add or remove it from the pool without any adverse effect on the volumes and their data for the pool, other than to impact the read-access performance.

You can create a read-cache disk group for a virtual pool under the following conditions:

- the system includes available SSDs, and
- the pool does not contain a virtual disk group comprised of SSDs (known as a Performance tier)


To increase the size of read cache within a pool, you must remove the read-cache disk group and then add a larger read-cache disk group. It is possible to have a read-cache disk group that consists of one or two disks. For more information on read cache, see [SSD read cache](#).

RAID levels

The controllers enable you to set up and manage disk groups, the storage for which may be spread across multiple disks.

This is accomplished through firmware resident in the controller. RAID refers to disk groups in which part of the storage capacity may be used to achieve fault tolerance by storing redundant data. The redundant data enables the system to reconstruct data if a disk in the disk group fails.

For a description of the ADAPT data protection level, see [ADAPT](#).

 **NOTE:** Choosing the right RAID level for your application improves performance.

In the PowerVault Manager you can create ADAPT, RAID-1, RAID-5, RAID-6, and RAID-10 disk groups. To create an NRAID or RAID-0 (linear only) disk group, you must use the `add disk-group` CLI command as described in the CLI Reference Guide.

The following tables:

- Provide examples of appropriate RAID levels for different applications.
- Compare the features of different RAID levels.
- Describe the expansion capability for different RAID levels (linear disk groups).
- Suggest the number of disks to select for different RAID levels (virtual disk groups).

Table 5. Example applications and RAID levels

Application	RAID level
Testing multiple operating systems or software development (where redundancy is not an issue)	NRAID
Fast temporary storage or scratch disks for graphics, page layout, and image rendering	0
Workgroup servers	1 or 10
Network operating system, databases, high availability applications, workgroup servers	5
Mission-critical environments that demand high availability and use large sequential workloads	6
Provides flexible storage and fast rebuilds. Well-suited for most workloads other than those using very few disks, or requiring a high number of writes	ADAPT

Table 6. RAID level comparison

RAID level	Min. disks	Description	Strengths	Weaknesses
NRAID	1	Non-RAID, nonstriped mapping to a single disk	Ability to use a single disk to store additional data	Not protected, lower performance (not striped)
0	2	Data striping without redundancy	Highest performance	No data protection: if one disk fails all data is lost
1	2	Disk mirroring	Very high performance and data protection; minimal penalty on write performance; protects against single disk failure	High redundancy cost overhead: because all data is duplicated, twice the storage capacity is required
5	3	Block-level data striping with distributed parity	Best cost/performance for transaction-oriented networks; very high performance and data protection; supports multiple simultaneous reads and writes; can also be optimized for large, sequential requests; protects against single disk failure	Write performance is slower than RAID 0 or RAID 1
6	4	Block-level data striping with double distributed parity	Best suited for large sequential workloads; non-sequential read and sequential read/write performance is comparable to RAID 5; protects against dual disk failure	Higher redundancy cost than RAID 5 because the parity overhead is twice that of RAID 5; not well-suited for transaction-oriented network applications; non-sequential write performance is slower than RAID 5
10 (1+0)	4	Stripes data across multiple RAID-1 subgroups	Highest performance and data protection (protects against multiple disk failures)	High redundancy cost overhead: because all data is duplicated, twice the storage capacity is required; requires minimum of four disks
ADAPT	12	Distributed erasure coding with dual disk failure protection supports 16+2 or 8+2 stripe width	Very fast rebuilds, no spare disks (built-in spare capacity), large storage pools, simplified initial deployment and expansion	Requires minimum of 12 disks

Table 7. Number of disks per RAID level to optimize virtual disk group performance

RAID level	Number of disks (data and parity)
1	2 total (no parity)
5	3 total (2 data disks, 1 parity disk); 5 total (4 data disks, 1 parity disk); 9 total (8 data disks, 1 parity disk)
6	4 total (2 data disks, 2 parity disks); 6 total (4 data disks, 2 parity disks); 10 total (8 data disks, 2 parity disks)
10	4–16 total
ADAPT	12–128 total

i NOTE: RAID 5 and RAID 6 disk groups, which have parity disks, should be created using the "power of 2" method to align properly with virtual pages. Failure to follow this method can result in significant degradation of sequential write performance.

RAID 5 disk groups should be created using 3, 5, or 9 disks. RAID 6 disk groups should be created using 4, 6, or 10 disks.

Table 8. Linear disk group expansion by RAID level

RAID level	Expansion capability	Maximum disks
NRAID	Cannot expand.	1

Table 8. Linear disk group expansion by RAID level (continued)

RAID level	Expansion capability	Maximum disks
0, 5, 6	You can add from 1 to 4 disks at a time.	16
1	Cannot expand.	2
10	You can add 2 or 4 disks at a time.	16
ADAPT	You can add up to 68 disks at a time.	128

ADAPT

ADAPT is a RAID-based data protection level that:

- Maximizes flexibility
- Provides built in spare capacity
- Optimizes performance
- Allows for very fast rebuilds, large storage pools, and simplified expansion

If a disk fails in an ADAPT disk group, and the failed disk is replaced with a new disk in the same slot, the replacement disk will be added to the disk group automatically. All disks in the ADAPT disk group must be the same type (enterprise SAS, for example), but can have different capacities, provided the range of difference does not exceed a factor of two. For example, mixing a 600 GB disk and a 1.2 TB disk is acceptable; but mixing a 6 TB disk and a 16 TB disk could prove problematic. It is conceivable that a sizeable difference between mixed disk capacities (ratio greater than two) could prevent consuming space on disks due to insufficient distributed space required to support striping.

NOTE: Do not mix disks if the ratio of the largest disk to the smallest disk is greater than two.

All disks in an ADAPT disk group are used to hold user data, but not all disks will be used by each page of data. To increase fault tolerance, any available capacity on disks can be allocated as spare for reconstruction purposes. When new data is added, new disks are added, or the system recognizes that data is not distributed across disks in a balanced way, the system moves the data to maintain balance across the disk group.

Spare disks are not used by ADAPT disk groups since the RAID design provides built-in spare capacity that is spread across all disks in the disk group. In the case of a disk failure, data will be redistributed to many disks in the disk group, allowing for quick rebuilds and minimal disruption to I/O.

The system will automatically default to a target spare capacity that is the sum of the largest two disks in the ADAPT disk group, which is large enough to fully recover fault tolerance after loss of any two disks in the disk group. The actual spare capacity value can change depending on the current available spare capacity in the disk group. Spare capacity is determined by the system as disks are added to a disk group, or when disk groups are created, expanded, or rebalanced.

NOTE: If a disk fails in an ADAPT disk group and is replaced by a new disk in the same slot as the failed disk, the disk group automatically incorporates the replacement disk into the disk group.

NOTE: For information about manually setting spare size, see the `add disk-group` command in the CLI Reference Guide. The `spare-capacity` parameter enables setting of the target spare capacity for an ADAPT disk group.

ADAPT disk groups can be expanded to either replenish current target spare capacity or to increase usable capacity. You can expand an ADAPT disk group from the **Maintenance > Storage** panel.

A system using ADAPT disk groups cannot be downgraded to a system that does not support ADAPT.

Disk group utilities


This section provides information about disk group utilities.

Disk group expansion

You can expand the capacity of an existing linear disk group. The RAID protection level determines whether the disk group can be expanded and the maximum number of disks the disk group can have. This task cannot be performed on an NRAID or RAID-1 disk group. Host I/O to the disk group can continue while the expansion proceeds.

Virtual disk-group expansion is only available with the ADAPT protection level. An ADAPT virtual disk group can consist of 12–128 disks of the same type. Host I/O to the disk group can continue while the expansion proceeds. This task cannot be performed on a non-ADAPT virtual disk group.

When expanding a disk group, all disks in the group must be the same type (enterprise SAS, for example). Disk groups support a mix of 512n and 512e disks. However, for best performance, all disks should use the same sector format. For more information about disk groups, see [Disk groups](#).

 **NOTE:** To expand a virtual pool, add a disk group as described in [Disk groups](#), or expand an ADAPT virtual disk group.

ADAPT disk groups are expanded when disks are added to the group. The controller determines how the additional disks are used, either to replenish spare capacity to equal target capacity, to increase usable capacity, or both.

For ADAPT disk groups, expansion is very fast. The controller invokes rebalance to redistribute spare capacity evenly across all disk members of the group to allow uniformly distributed usable capacity. Due to the possible need to rebalance, and to maintain fault tolerance and target spare capacity, any new usable capacity may not be immediately available. Monitor the Activity panel for progress on these activities, and check for updated usable capacity when the activities are complete. When set to the default spare capacity, the system will try to replenish spare capacity to be the sum of the largest two disks in the group.

In the **Maintenance > Storage** panel, locate the ADAPT disk group to expand, access its slide-over panel, click **Expand Disk Group**, and follow the on-screen directions.

Expanding non-ADAPT linear disk groups

Before expanding non-ADAPT disk groups, back up the data in the disk group so that if you need to stop expansion and delete the disk group, you can move the data into a new, larger disk group.

In the **Maintenance > Storage** panel, locate the non-ADAPT disk group to expand, display its slide-over panel, click **Expand Disk Group**, and follow the on-screen directions.

To cancel disk group expansion, click the  icon

Disk-group scrub

The disk-group scrub utility analyzes specified disk groups to find and fix errors. It can find media errors for any protection level and for a read cache disk group. By default, the utility is enabled to run periodically.

The disk-group scrub utility acts on all disks in the disk group, but not leftover disks. The disk-group scrub utility performs these tasks:

- Checks redundancy data (parity) and corrects it for protection levels 5, 6, and ADAPT .
- Finds, but does not fix, mirror mismatches for protection levels 1 and 10.
 - The system reads both copies of mirror data to find any mismatches.
- Finds and fixes media errors for all redundant protection levels.
 - Media errors occur when the system cannot read one of the copies of mirror data, due to a disk error such as an Unrecoverable Read Error (URE).
 - (RAID-1 and RAID-10). Verify that all blocks are readable (NRAID and RAID-0).

A disk-group scrub can last well over an hour, depending on disk-group size, utility priority, and the amount of I/O activity. You can use a disk group while it is being scrubbed. While the scrub is running, you can monitor progress and cancel if necessary. When the scrub is complete, event 207 is logged.

To run the scrub utility, see [Scrubbing a disk group](#).

NOTE: CLI commands related to disk-group scrub:

- You can set the disk scrub duration (hours) by using the `scrub-duration-goal` parameter of the `set disk-group` CLI command
- You can query disk group properties related to the media scrub utility by using the `scrub-status` parameter of the `show disk-groups` CLI command.

See the CLI Reference Guide for more information.

SSDs

The use of SSDs (solid-state drives) can greatly enhance the performance of a system. Since the SSDs do not have moving parts, data that is random in nature can be accessed much faster.

You can use SSDs for virtual disk groups in a system that contains only SSDs, or in a system that uses SSDs in combination with spinning disks. When combined with virtual disk groups that consist of other classes of disks, improved read and write performance is possible through automated tiered storage. Alternatively, you can use one or two SSDs in read-cache disk groups to increase read performance for pools without a Performance tier. The application workload of a system determines the percentage of SSDs of the total disk capacity that is needed for best performance.

The use of SSDs may also greatly enhance the performance of a linear system for the same reasons they benefit a virtual system. You can use SSDs in combination with spinning disks. The application workload of a system determines the percentage of SSDs of the total disk capacity that is needed for best performance.

Follow these recommendations for using SSDs and spinning disks:

- If the first disk group is provisioned with SSDs, then the system will expect to be provisioned as an all-flash array and allow only SSDs to be used in all other disk groups.
- If the first disk group is provisioned with spinning disks, then the system can only be provisioned to use spinning disks in virtual disk groups and use SSDs as read cache.
- If the first disk group is provisioned with spinning disks or SSDs, then the system can be provisioned to use spinning disks in virtual disk groups and use SSDs either in virtual disk groups or as read cache.

Related information

[Read-cache disk groups](#)

[Gauging the percentage of life remaining for SSDs](#)

[SSD management](#)

[SSD read cache](#)

Gauging the percentage of life remaining for SSDs

An SSD can be written and erased a limited number of times.

Through the SSD Life Left disk property, you can gauge the percentage of disk life remaining. This value is polled every 5 minutes. When the value decreases to 20%, an event is logged with Informational severity. This event is logged again with Warning severity when the value decreases to 5%, 2% or 1%, and 0%. If a disk crosses more than one percentage threshold during a polling period, only the lowest percentage will be reported. When the value decreases to 0%, the integrity of the data is not guaranteed. To prevent data integrity issues, replace the SSD when the value decreases to 5% of life remaining.

Under **Maintenance > Hardware** within the PowerVault Manager, select the SSD. See the SSD Life Left label under Disk Information to view the percentage of SSD life remaining. To view the percentage of SSD life remaining using the CLI, enter the `show disks` CLI command with the `detail` parameter as described in the CLI Reference Guide.


Viewing I/O workload activity

The PowerVault Manager provides the I/O workload graph to help you configure the target SSD size to get the best performance for your workload. This graph presents a heat map of system activity.

The I/O workload graph shows the relationship between the workload and the amount of storage capacity used. This data reveals how much capacity is frequently accessed over time ("hot"). Under most workloads, the graph is a good indicator of data that the tiering algorithm would have put on SSDs if sufficient SSD capacity existed.


You can use this information to determine how system performance can benefit from implementing a tier of fast SSDs, instead of slower spinning disks for some or all of that capacity. Users often see the greatest performance benefits when the SSD tier is sized to handle 80% or more of the I/O workload.

Calculations are based on user-specified settings and up to eight days of usage data captured by the system.

 **NOTE:** The suggested capacities might not apply to heavily streaming workloads.

Access the I/O Workload graph from the Dashboard:

1. Select the Capacity panel slide-over.
2. In the pool detail area for the appropriate pool, select **View I/O Workload**.

 **NOTE:** Because the I/O Workload graph is accessed through a specific pool (A or B) on the Capacity panel, this information is available only if the system is provisioned for virtual storage. Pools are not displayed for linear storage.

You can set the following options

- **Values**—Select whether to base the calculations on the peak values saved in the usage data or the average values. For calculations, the pool is divided into equal "bins" of LBAs. Each sample contains readings for all bins. Multiple samples are taken per day. To calculate average, the sum of the readings of a bin are divided by the number of samples. To calculate peak, the largest bin value from the collection of samples is taken. This process leaves one value for each bin, whether average or peak is selected. From there, workload calculations are made using the bin numbers as input.
- **Show**—Select whether to limit the data used for calculations to read I/Os only, write I/Os only, or the combined total of read and write I/Os
- **Workload**—Select from one to three workload calculations to display. The default calculations are based on low, mid, and high percentages of capacity: 50%, 80%, and 100%. In place of 50%, you can enter a custom percentage, which must be a whole number.

Reading the graph

The graph contains a line that reflects the capacity and a line plot for each selected workload.

- When graphed elements are above the SSD capacity line (or if there are no SSDs), data is spread over more capacity in the total system than could be serviced by the SSD capacity. The graph can give you a target SSD size to consider using.
- When graphed elements are below the SSD capacity line, there is adequate SSD capacity for hot data and you're receiving good value from your SSDs.

Interpreting this graph requires you to balance your expectations of cost versus performance. For example, you might be willing to have a couple of days where peak usage is far above the capacity line because it is acceptable to have slower performance during these times, given the cost. Or you could design your system to perform well during those times so the system has good I/O performance at all times.

Cumulative heat map data can also be useful in support case analysis. For information, see [Downloading system heat map data](#).

SSD management

SSDs use multiple algorithms to manage SSD endurance features. These include wear leveling to prolong service life, support for Unmap commands, and over-provisioning to minimize write amplification. SSDs use data retention algorithms to monitor and mitigate cell level decay.

Wear leveling

Wear leveling is a technique for prolonging the service life of some kinds of erasable computer storage media, such as the flash memory used in SSDs. It attempts to ensure that all flash cells are written to or exercised as evenly as possible to avoid any hot spots where some cells are used up faster than other locations. There are several different wear leveling mechanisms used in flash memory systems, each with different levels of success.

Vendors have different algorithms to achieve optimum wear leveling. Wear leveling management occurs internal to the SSD. The SSD automatically manages wear leveling, which does not require any user interaction.

Overprovisioning

The write amplification factor of an SSD is defined as the ratio of the amount of data actually written by the SSD to the amount of host/user data requested to be written. The write amplification factor affects wear-leveling calculations and is influenced by the characteristics of data written to and read from SSDs.

Data written in sequential LBAs aligned on 4KB boundaries results in the best write amplification factor. The worst write amplification factor typically occurs for randomly written LBAs of transfer sizes less than 4KB, originating on LBAs not on 4KB boundaries. Try to align your data on 4KB boundaries.

UNMAP command

The `UNMAP` command in the SCSI command set allows an operating system to inform an SSD of the blocks of data that are no longer considered in use and can be wiped internally.

Data retention

Data retention is another major characteristic of SSDs that all SSD algorithms take into account while running.

While powered up, the data retention of SSD cells is monitored and rewritten if the cell levels decay to an unexpected level. Data retention when the drive is powered off is affected by Program and Erase (PE) cycles and the temperature of the drive when stored.

Drive writes per day (DWPD)

Disk vendors rate SSD endurance by how many writes can occur over the lifetime of an SSD. As lower-cost SSDs that support fewer drive writes per day become available, the cost/benefit analysis regarding which SSDs to use is highly dependent on your applications and I/O workload, together with the ratio of SSDs to conventional disks.

Since the storage system tiering algorithm automatically moves "hot" data to SSDs and less-used "cool" data to conventional disks, applications and environments that require mission-critical movement of frequently accessed "hot" data might dictate a higher ratio of SSDs to conventional disks.

Because data is characterized every five seconds and moved to the appropriate storage device, no fixed rule is used to determine which SSDs are used. For this reason, using SSDs with the same DWPD values is advised.

SSD read cache

Unlike tiering, where a single copy of specific blocks of data resides in either spinning disks or SSDs, the read cache feature uses one SSD read cache disk group per pool as a read cache for "hot" pages only. If a virtual pool contains a flash tier, then read-cache is not allowed.

Each read-cache disk group consists of one or two SSDs with a maximum usable capacity of 4TB. A separate copy of the data is also kept in spinning disks. Read cache contents are lost when a controller restart or failover occurs. This does not cause data loss or corruption, as the read cache only duplicates the content that exists in the fault-tolerant disk groups. Taken together, these attributes have several advantages:

- The performance cost of moving data to read-cache is lower than a full migration of data from a lower tier to a higher tier.
- Read-cache does is not fault tolerant, potentially lowering system cost.
- Controller read cache is effectively extended by two orders of magnitude, or more.
- When a read-cache group consists of one SSD, it automatically uses NRAID. When a read-cache group consists of two SSDs, it automatically uses RAID 0.

For more information, see [SSDs](#).

Spares

Spare disks are unused disks in your system that automatically replace a failed disk, restoring fault tolerance to disk groups in the system. Designate spares from the **Maintenance > Storage** panel or by using the `add spares` CLI command. For information about this command, see the CLI Reference Guide.

Types of spares include:

- **Dedicated spare**—Reserved for use by a specific linear disk group to replace a failed disk. Most secure way to provide spares for disk groups, but it is expensive to reserve a spare for each disk group.
- **Global spare**—Reserved for use by any fault-tolerant disk group to replace a failed disk.
- **Dynamic spare**—Available compatible disk that is automatically assigned to replace a failed disk in a fault-tolerant disk group.

NOTE: ADAPT disk groups do not use global spares or dynamic spares. For information on how ADAPT disk groups manage sparing, see [ADAPT](#).

A controller automatically reconstructs a fault-tolerant disk group (RAID 1, 5, 6, 10) when one or more of its disks fails and a compatible spare disk is available. A disk is compatible if it has enough capacity to replace the failed disk and is the same type (enterprise SAS, for example). If the disks in the system are FDE-capable and the system is secure, spares must also be FDE-capable.

NOTE: Sufficient disks must remain in the disk group so that reconstruction is possible.

When a disk fails, the system looks for a dedicated spare first. If it does not find a dedicated spare, it looks for a global spare. If it does not find a compatible global spare and the dynamic spares option is enabled, it takes any available compatible disk. If no compatible disk is available, reconstruction cannot start.

On a 5U84 enclosure, the system first looks for spares in the same drawer as the failed disk, then in another drawer in the same enclosure, and then in another enclosure within the same system.

NOTE: A best practice is to designate spares for use if disks fail. Dedicating spares to disk groups is the most secure method, but it is also expensive to reserve spares for each disk group. Alternatively, you can enable dynamic spares or assign global spares.

Pools

Virtual and linear storage systems both use pools. A pool is an aggregation of one or more disk groups that serves as a container for volumes.

A disk group is a group of disks of the same type, using a specific RAID level. For virtual pools, when volumes are added to the pool, the data is distributed across the disk groups in the pool. For linear pools, which have only one disk-group, the data is distributed linearly across the pool. If the owning controller fails, the partner controller assumes temporary ownership of the pool and resources owned by the failed controller. If a fault-tolerant cabling configuration, with appropriate mapping, is used to connect the controllers to hosts, LUNs for both controllers are accessible through the partner controller so I/O to volumes can continue without interruption.


NOTE: Corresponding to the two storage methods, there are both virtual and linear pools and disk groups. There is another type of disk group—the read-cache disk group—which is also related to virtual storage. Read-cache disk groups consist of SSDs. If your system does not use SSDs, you will not be able to create read-cache disk groups.

In both virtual and linear storage, if the owning controller fails, the partner controller assumes temporary ownership of the pool and resources owned by the failed controller. If a fault-tolerant cabling configuration, with appropriate mapping, is used to connect the controllers to hosts, LUNs for both controllers are accessible through the partner controller so I/O to volumes can continue without interruption.

You can provision disks into disk groups. For information about how provisioning disks works, see [Adding a disk group to a pool](#).

Virtual pools and disk groups

The volumes within a virtual pool are allocated virtually (separated into fixed size pages, with each page allocated according to the algorithm within the tiers and disk groups in the pool) and thinly (meaning that they initially exist as an entity, but don't have any physical storage allocated to them). They are also allocated on-demand (as data is written to a page, it is allocated).

 **NOTE:** The physical capacity limit for a virtual pool is 4 PiB.


You can remove one or more disk groups, but not all, from a virtual pool without losing data if there is enough space available in the remaining disk groups to move the data into. When the last disk group is removed, the pool ceases to exist, and will be deleted from the system automatically. Alternatively, the entire pool can be deleted, which automatically deletes all volumes and disk groups residing on that pool. Deleting a pool that contains data will cause permanent data loss.


If a system has at least two SSDs, each virtual pool can also have a read-cache disk group. Unlike the other disk group types, read-cache disk groups are used internally by the system to improve read performance and do not increase the available capacity of the pool. See [Read-cache disk groups](#) for more information.

Resolving a pool conflict caused by inserting a foreign disk group

If you insert a virtual disk group, a disk that contains a virtual disk group, or a disk that was actively part of a virtual disk group from one system into another system, the latter system will attempt to create a virtual pool for that disk group.

If that system already has a virtual pool with the same name, the pool for the inserted disk group will be offline. For example, if NewSystem has pool A and you insert a disk group that came from OldSystem pool A, the disk group imported from OldSystem pool A will be offline.

 **NOTE:** You cannot access the two sets of pool A data (OldSystem and NewSystem) concurrently on a single storage system.

 **NOTE:** If you are unable to find a pool with a duplicate name, or are unsure of how to safely proceed, download logs from the system and contact technical support for assistance with offline disk groups.

Linear pools and disk groups


Each time that the system adds a linear disk group, it also creates a corresponding pool for the disk group. Once a linear disk group and pool exists, volumes can be added to the pool. The volumes within a linear pool are allocated in a linear/sequential way, such that the disk blocks are sequentially stored on the disk group.

Linear storage maps logical host requests directly to physical storage. In some cases the mapping is 1-to-1, while in most cases the mapping is across groups of physical storage devices, or slices of them.


Changing pool settings

Each virtual pool has three thresholds for page allocation as a percentage of pool capacity. You can set the low and middle thresholds. The high threshold is automatically calculated based on the available capacity of the pool minus 200 GB of reserved space.

You can view and change settings that govern the operation of each virtual pool from the pools panel (**Maintenance >**

Storage): To see information about disk groups in a pool, expand the row. To change pool settings, click the  icon in the pool row. Options include:

- **Low Threshold**—When this percentage of virtual pool capacity has been used, the system will generate an alert and informational event 462 to notify the administrator. This value must be less than or equal to the Middle Threshold value. The default is 50%.
- **Middle Threshold**—When this percentage of virtual pool capacity has been used, the system will generate an alert and event 462 to notify the administrator to add capacity to the pool. This value must be between the Low Threshold and High Threshold values. The default is 75%. If the pool is not overcommitted, the alert will have Informational severity. If the pool is overcommitted, the alert will have Warning severity.
- **High Threshold**—When this percentage of virtual pool capacity has been used, the system will generate an alert and event 462 to alert the administrator to add capacity to the pool. This value is automatically calculated based on the available capacity of the pool minus 200 GB of reserved space. If the pool is not overcommitted, the alert will have Informational severity. If the pool is overcommitted, the alert will have Warning severity and the system will use write-through cache mode until virtual pool usage drops back below this threshold.
- **Pool Overcommit**—This check box controls whether overcommitting is enabled, and whether storage-pool capacity may exceed the physical capacity of disks in the system. For information about overcommitting, [Overcommitting volumes](#).

 **NOTE:** For more information about events, see the Event History panel (**Maintenance > Support > Event History**).

NOTE: If your system has a replication set, the pool might be unexpectedly overcommitted because of the size of the internal snapshots of the replication set.

NOTE: If the pool is overcommitted and has exceeded its high threshold, its health will show as degraded in the Storage panel (**Maintenance > Storage**). If you try to disable overcommitment and the total space allocated to thin-provisioned volumes exceeds the physical capacity of their pool, an error will state that there is insufficient free disk space to complete the operation and overcommitment will remain enabled.

To check if the pool is overcommitted, go to **Maintenance > Storage**, then expand the pool row. If the Pool Overcommitted value is **True**, the pool is overcommitted. If the value is **False**, the pool is not overcommitted.

Volumes and volume groups

A volume is a logical subdivision of a virtual or linear pool and can be attached to hosts.

An attached volume provides addressable storage to a host (for example, a file system partition you create with your operating system or third-party tools). For more information, see [Attaching volumes to hosts](#). For virtual pools, when volumes are added to a pool the data is distributed across the disk groups in the pool. For linear pools, which can only have one disk group per pool, volumes are also added to the pool, which contains the volume data.

Virtual volumes

Virtual volumes make use of a method of storing user data in virtualized pages. These pages may be spread throughout the underlying physical storage and allocated on demand. Virtualized storage therefore has a dynamic mapping between logical and physical blocks.

A maximum of 1024 virtual volumes can exist per system.

Volume groups

You can group a maximum of 1024 volumes (standard volumes, snapshots, or both) into a volume group. Doing so enables you to perform mapping operations for all volumes in a group at once, instead of for each volume individually.

NOTE: Volume groups are only supported through the CLI.

A volume can be a member of only one group. All volumes in a group must be in the same virtual pool. A volume group cannot have the same name as another volume group, but can have the same name as any volume. A maximum of 256 volume groups can exist per system. If a volume group is being replicated, the maximum number of volumes that can exist in the group is 16.

Linear volumes

Linear volumes make use of a method of storing user data in sequential fully allocated physical blocks. Mapping between the logical data presented to hosts and the physical location where it is stored is fixed, or static. These blocks have a fixed (static) mapping between the logical data presented to hosts and the physical location where it is stored.

Volume cache options

You can set options that optimize reads and writes performed for each volume. It is recommended that you use the default settings.

For more information, see the following topics:

- [Using write-back or write-through caching](#)
- [Cache optimization mode](#)
- [Optimizing read-ahead caching](#)

Using write-back or write-through caching

You can change the cache policy setting of a volume in a dual-controller system.

CAUTION: In a dual-controller system, only disable write-back caching if you fully understand how the host operating system, application, and adapter move data. System performance may be impacted if the caching profile is used incorrectly.

Write-back is a cache-writing strategy in which the controller receives the data to be written to disks, stores it in the memory buffer, and immediately sends the host operating system a signal that the write operation is complete, without waiting until the data is written to the disk. Write-back cache is saved to nonvolatile storage in the event of a power loss. Data in the cache is mirrored from one controller module cache to the other. If a controller fails, the remaining controller completes the write operation to the disks. Write-back cache improves the performance of write operations and the throughput of the controller.

In write-through caching mode, the controller writes the data to the disks before signaling the host operating system that the process is complete. Write-through cache has lower write throughput performance than write-back, but it is the safer strategy, with minimum risk of data loss on power failure. Write-through cache does not require mirroring the write data because the data is written to the disk before signaling that the command is complete. You can set conditions that cause the controller to change from write-back caching to write-through caching.

In both caching strategies, active/active failover of the controllers is enabled.

You can enable and disable the write-back cache for each volume:

- In a dual-controller system, with both controllers running, volume write-back cache is enabled by default. Because controller cache is backed by supercapacitor technology, if the system loses power, data is not lost. For dual-controller systems, write-back is the preferred setting.
- In a single-controller system, or if a controller fails in a dual-controller system, write-through cache is enabled by default.

NOTE: The best practice for a fault-tolerant configuration is to use write-back caching.

Cache optimization mode

CAUTION: Changing the cache optimization setting while I/O is active can cause data corruption or loss. Before changing this setting, quiesce I/O from all initiators.

You can change the cache optimization mode to one of the following modes of operation.

- **standard**—In this mode of operation, the controller sets volume cache parameters to address both sequential I/O and random I/O tasks. This optimization is the choice for most workloads. In this mode, the cache is kept coherent with the partner controller. This mode provides high performance and high redundancy, and it is the default.
- **cache-hit**—This controller cache mode of operation is optimized for workloads that are localized, that is, a substantial percentage of all accesses are hits in the controller cache. In this mode, the cache is kept coherent with the partner controller.
- **standard-atomic-write**—This controller cache mode of operation adds the atomic-write feature to the standard mode. Atomic Write guarantees that when a Host Write is interrupted for any reason, the data returned on subsequent read will be either the "old" data accepted prior to the interrupted Write or the "new" data from the interrupted write.
- **cache-hit-atomic-write**—This controller cache mode of operation adds the atomic-write feature to the cache-hit mode. Atomic Write guarantees that when a Host Write is interrupted for any reason, the data returned on subsequent read will be either the "old" data accepted prior to the interrupted Write or the "new" data from the interrupted write.

Optimizing read-ahead caching

You can optimize a volume for sequential reads or streaming data by changing its read-ahead cache settings.

CAUTION: Only change read-ahead cache settings if you fully understand how the host operating system, application, and adapter move data so that you can adjust the settings accordingly.

You can change the amount of data read in advance. Increasing the read-ahead cache size can greatly improve performance for multiple sequential read streams.

- The **Adaptive** option works well for most applications: it enables adaptive read-ahead, which allows the controller to dynamically calculate the optimum read-ahead size for the current workload.
- The **Stripe** option sets the read-ahead size to one stripe. The controllers treat NRAID and RAID-1 disk groups internally as if they have a stripe size of 512 KB, even though they are not striped.

- Specific size options let you select an amount of data for all accesses. Options include 512 KB, 1 MB, 2 MB, 4 MB, 8 MB, 16 MB, 32 MB.
- The **Disabled** option turns off read-ahead cache. This is useful if the host is triggering read ahead for what are random accesses. This can happen if the host breaks up the random I/O into two smaller reads, triggering read ahead.

Overcommitting volumes

Overcommitting is a virtual storage feature that allows a system administrator to overcommit physical storage resources. This allows the host system to operate as though it has more storage available than is actually allocated to it. When physical resources fill up, the administrator can add physical storage by adding additional disk groups or expanding an ADAPT disk group.


Paging is required to eliminate the lack of flexibility associated with linear mapping. Linear mapping limits the ability to easily expand the physical storage behind the over-provisioned volume. Paged mapping allows physical resources to be disparate and noncontiguous, making it much easier to add storage on the fly.

About automated tiered storage

Automated tiered storage is a virtual storage feature that automatically moves data residing in one class of disks to a more appropriate class of disks based on data access patterns, with no manual configuration necessary. Automated tiered storage operates as follows:

- Frequently accessed "hot" data can move to disks with higher performance.
- Infrequently accessed "cool" data can move to disks with lower performance and lower costs.
- Each virtual disk group, depending on the type of disks it uses, is automatically assigned to one of the following tiers:
 - **Performance**—This highest tier uses SSDs, which provide the best performance but also the highest cost. For more information on SSDs, see [SSDs](#).
 - **Standard**—This middle tier uses enterprise-class spinning SAS disks, which provide good performance with mid-level cost and capacity.
 - **Archive**—This lowest tier uses midline spinning SAS disks, which provide the lowest performance with the lowest cost and highest capacity.

When the status of a disk group in the Performance tier becomes CRIT (critical), the system will automatically drain data from that disk group to disk groups using spinning disks in other tiers, provided that those disk groups contain the data on the degraded disk group. This occurs because similar wear across the SSDs is likely, so more failures may be imminent.

 **NOTE:** If the system uses only SSDs, data is moved to another fault tolerant SSD tier that has free capacity.

Automated tiered storage rebalancing happens when adding or removing a disk group in a system.

Volume tier affinity

The volume tier affinity feature enables tuning the tier-migration algorithm for a virtual volume when creating or modifying the volume so that the volume data automatically moves to a specific tier, if possible. If space is not available in the preferred tier, another tier will be used.

The three volume tier affinity settings are:

- **No Affinity** —This setting uses the highest available performing tiers first and only uses the Archive tier when space is exhausted in the other tiers. Volume data swaps into higher performing tiers based on the frequency of access and tier space availability.
- **Archive**—This setting prioritizes the volume data to the lowest tier of service. Volume data can move to higher performing tiers based on the frequency of access and available space in the tiers.
- **Performance**—This setting prioritizes volume data to the higher performing tiers. If no space is available, lower performing tier space is used. Performance affinity volume data swaps into higher tiers based upon frequency of access or when space is made available.

Initiators, hosts, and host groups

An initiator represents an external port to which the storage system is connected. The external port may be a port in an I/O adapter (such as HBA) in a server.

For ease of management, you can group 1 to 128 initiators that represent a server into a host. A host is a user-defined object that represents a server to which the storage system is attached, and is used to define a mapping relationship to storage.

Further, you can also group 1 to 256 hosts into a host group. A host group is a user-defined set of hosts. Doing so enables you to attach all grouped initiators in a host, or all initiators and hosts in a group, instead of for each initiator or host individually.

The controllers automatically discover initiators that have sent a SCSI `INQUIRY` command or a `REPORT LUNS` command to the storage system, which typically happens when a host boots up or rescans for devices. When the command is received, the system saves the initiator ID. You can also manually create entries for initiators as described in the CLI Reference Guide by setting a nickname to a specified unique ID. For example, you might want to define an initiator before a controller port is physically connected through a switch to a server.

In the PowerVault Manager, you must assign a nickname to an initiator in order for it to be added to a host. An initiator can be a member of only one host. A host can be a member of only one group. A host cannot have the same name as another host, but can have the same name as any initiator. A host group cannot have the same name as another host group, but can have the same name as any host. A maximum of 32 host groups can exist. Once you have created a host, you can edit the profile specific to the operating system for that initiator.

CHAP

A storage system with iSCSI ports can be protected from unauthorized access via iSCSI by enabling Challenge Handshake Authentication Protocol (CHAP).

CHAP authentication occurs during an attempt by a host to log in to the system. This authentication requires an identifier for the host and a shared secret between the host and the system. Optionally, the storage system can also be required to authenticate itself to the host. This is called mutual CHAP. You are prompted to optionally configure CHAP settings during the onboarding process. Once onboarding is complete, you can enable or disable CHAP and create new CHAP records from the **Settings > iSCSI** panel. Steps involved to enable CHAP include:

- Decide on host node names (identifiers) and secrets. The host node name is its IQN. A secret must have 12-16 characters, and include spaces and printable UTF-8 characters except: " or <

This authentication requires an identifier for the host and a shared secret between the host and the system. The CHAP secret is a text string that is known to both the initiator and the storage array before they negotiate a communication link. Mutual CHAP authenticates the target to the initiator. Without mutual CHAP, only the initiator is authenticated to the target.

- Define CHAP records in the storage system.
- Enable CHAP on the storage system (during onboarding or from the **Settings > iSCSI > Configuration** panel). Note that this applies to all iSCSI hosts, in order to avoid security exposures. Any current host connections will be terminated when CHAP is enabled and will need to be re-established using a CHAP login.
- Define a CHAP record for the host iSCSI initiator on the host.
- Establish a new connection to the storage system using CHAP. The host should be displayable by the storage system, as well as the ports through which connections were made.

 **CAUTION: Changing iSCSI configuration settings after onboarding can invalidate CHAP settings. This could disrupt connectivity between the host and the storage system.**

If it becomes necessary to add more hosts after CHAP is enabled, additional CHAP node names and secrets can be added. If a host attempts to log in to the storage system, it will become visible to the system, even if the full login is not successful due to incompatible CHAP definitions. This information may be useful in configuring CHAP entries for new hosts, and becomes visible when an iSCSI discovery session is established because the storage system does not require discovery sessions to be authenticated. CHAP authentication must succeed for normal sessions to access LUNs from the storage array. To use CHAP between peers in a replication set, see [CHAP and replication](#).

Host ports

Controller enclosures support FC, iSCSI, and SAS host interface protocols. FC and SAS controller host-port settings are not configurable in the PowerVault Manager.


iSCSI controller host-port settings should be configured, preferably, during onboarding, in order to enable the system to communicate with iSCSI hosts. This process includes choosing the systems iSCSI network configuration type (either IPv4 or IPv6), entering the IP address of at least one host port on each controller, and providing the netmask and gateway for assigned port IP addresses. When initial iSCSI configuration is complete, you can view and change host port settings.

 **CAUTION:** Changing host port settings while initiators are in use can disrupt host to LUN connections.


Attaching volumes to hosts


A volume must be attached to one or more hosts (or host groups) to enable them to access the volume.


You can attach a volume to hosts as part of creating the volume, or afterward. When attaching a volume you can choose whether to create new hosts, or to use existing hosts. For information about creating hosts, see [Attaching volumes to hosts](#). When an attachment is created, the system automatically assigns a unique LUN to the volume, sets default permission access to read-write, and sets port access to all ports. After an attachment is created, you can change the LUN, port access, and access permissions. Both controllers share a set of LUNs, and any available LUN can be assigned to a volume. The storage system uses Unified LUN Presentation (ULP), which can expose all LUNs through all host ports on both controllers. The interconnect information is managed by the controller firmware. ULP appears to the host as an active-active storage system where the host can choose any available path to access a LUN regardless of which controller owns the storage pool the volume resides on. With ULP, the controllers' operating/redundancy mode is shown as Active-Active ULP. ULP uses the T10 Technical Committee of INCITS Asymmetric Logical Unit Access (ALUA) extensions, in SPC-3, to negotiate paths with aware host systems. Unaware host systems see all paths as being equal.

 **NOTE:** LUN 0 is not used for SAS hosts.

The system also sets properties that specify whether the volume is attached to at least one host, whether the host was discovered, and whether the volume is accessible through redundant paths (through host ports in each controller module).

 **NOTE:** The secondary volume of a replication set cannot be attached to hosts. To enable such access, create a snapshot of the secondary volume and attach the snapshot to the host.

 **NOTE:** To avoid multiple hosts mounting the volume and causing data integrity issues, the host computer systems must be cooperatively managed, such as by using cluster software. If multiple hosts mount a volume without being cooperatively managed, volume data is at risk for data integrity failures.

 **CAUTION:** Volume attachment changes take effect immediately. Make changes to volumes when the volumes are not in use. Before changing a LUN, be sure to unmount the volume.

You can perform the following attachment actions:

- View information about hosts attached to a volume (**Provisioning > Volumes**)
- Attach volumes to hosts or host groups (**Provisioning > Volumes > Attach to Hosts**)
- Detach volumes from hosts or host groups (**Provisioning > Volumes > Detach from Hosts**)
- View information about volumes attached to a host (**Provisioning > Hosts**)
- Attach hosts to volumes (**Provisioning > Hosts > Attach to Volumes**)
- Detach hosts from volumes (**Provisioning > Hosts > Detach from Volumes**)

Operating with a single controller

An ME5 system with a single controller module does not offer a redundant configuration. If the controller fails, this leaves the system at risk for data unavailability.

A single-controller system defaults to write-through caching, in which data is written directly to disks rather than to nonvolatile storage. This setting cannot be changed in a single-controller system.

NOTE: If you are operating a system with a single controller, some functionality that is described in the documentation may be unavailable or not applicable to your system. For example, only one storage pool can exist and information about controller failover and recovery is not applicable.

Snapshots

The system can create snapshots of volumes up to the maximum number supported by your system. Snapshots provide data protection by enabling you to create and save source volume data states at the point in time when the snapshot was created. Snapshots can be created manually or you can schedule snapshot creation. After a snapshot has been created, the source volume can be expanded, but the snapshot volume cannot be expanded.

NOTE: Data protection and copy services are supported for virtual storage only; these features are not supported if the storage type is set to linear.

To view the maximum number of snapshots for your system, see [System configuration limits](#).

When you reach the maximum number of snapshots for your system, before you can create a new snapshot, you must delete an existing snapshot.

The system treats a snapshot like any other volume.

- Snapshots attach to hosts with read-write access by default. You can then change permissions to read-only access or no access, depending on the purpose of the snapshot.
- Snapshots use the rollback feature, which replaces the data of a source volume or snapshot with the data of a snapshot that was created from it.
- Snapshots also use the reset snapshot feature, which enables you to replace the data in a snapshot with the current data in the source volume. When you reset a snapshot, the snapshot name and mappings are not changed.

The `set snapshot-space` CLI command enables you to set the percent of the pool that can be used for snapshots (the snapshot space). Optionally, you can specify a limit policy to enact when the snapshot space reaches the percentage. You can set the policy to either notify you via events that the percentage has been reached (in which case the system continues to take snapshots, using the general pool space), or to notify you and trigger automatic deletion of snapshots. If automatic deletion is triggered, snapshots are deleted according to their configured retention priority.

Creating snapshots is a fast and efficient process that merely consists of pointing to the same data to which the source volume or snapshot points. (A snapshot takes up no space unless it is directly modified, or the source volume or source snapshot to which the snapshot refers to is modified.) Space does not have to be reserved for snapshots because all space in the pool is available for them, based upon the limit policy that is set. It is easy to take snapshots of snapshots and use them in the same way that you would use any volume. Since snapshots have the same structure as volumes, the system treats them the same way.

Snapshot trees

Because a snapshot can be the source of other snapshots, a single volume can be the progenitor of many levels of snapshots.

Originating from an original base volume, the levels of snapshots create a snapshot tree that can include up to 254 snapshots, each of which can also be thought of as a leaf of the tree. When snapshots in the tree are the source of additional snapshots, they create a new branch of the snapshot tree and are considered the parent snapshot of the child snapshots, which are the leaves of the branch.

The tree can contain snapshots that are identical to the volume or have content that has been later modified. Once the 254-snapshot limit has been reached, you cannot create additional snapshots of any item in the tree until you manually delete snapshots from the tree.

You can only delete snapshots that do not have any child snapshots. You can expand the base volume of a snapshot tree, but you cannot expand any snapshots in the tree.

Rollback and reset snapshot features

With the rollback feature, if the contents of the selected snapshot have changed since it was created, the modified contents will overwrite those of the source volume or snapshot during a rollback. Since snapshots are copies of a point in time, a modified snapshot cannot be reverted. If you want a virtual snapshot to provide the capability to revert the contents of the source volume or snapshot to when the snapshot was created, create a snapshot for this purpose and archive it so you do not change the contents. For more information see [Rolling back virtual volumes](#).

As an alternative to taking a new snapshot of a volume, you can replace the data in a standard snapshot with the current data in the source volume. The snapshot name and host attachments are not changed. For snapshots, the reset snapshot feature is supported for all snapshots in a tree hierarchy. However, a snapshot can only be reset to the immediate parent volume or snapshot from which it was created.

 **CAUTION:** To avoid data integrity issues, detach a snapshot from hosts before resetting the snapshot.

You can either reset a snapshot immediately, or schedule the reset. For more information, see [Resetting snapshots](#).

Copying volumes or snapshots

The volume copy feature (**Provisioning > Volumes > volume slide-over > Copy Volume**) enables you to copy a base volume or snapshot to a new volume. It creates a complete "physical" copy of a base volume or snapshot within a storage system. The copy is an exact duplicate of the source as it existed at the time the copy operation was initiated, consumes the same amount of space as the source, and is independent from an I/O perspective. In contrast, the snapshot feature creates a point-in-time "logical" copy of a volume, which remains dependent on the source volume.

The volume copy feature provides the following benefits:

- **Additional data protection:** An independent copy of a volume provides additional data protection against a complete source volume failure. If the source volume fails, the copy can be used to restore the volume to the point in time when the copy was created.
- **Non-disruptive use of production data:** With a mounted independent copy of the volume, resource contention and the potential performance impact on production volumes is mitigated. Data blocks between the source and the copied volumes are independent (versus shared with snapshots) so that I/O is to each set of blocks respectively. Application I/O transactions are not competing with each other when accessing the same data blocks.

The copy operation is performed directly from the source. This source data may change if modified data is to be included in the copy and there are snapshots attached and in use. You must detach the volume before copying it. It is highly recommended to unmount it on the host first, so that the host can flush its cache, before detaching it. The volume will not be available for read or write access until the copy is complete, at which time you can reattach the volume.


Reconstruction

If one or more disks fail in a disk group, adequate disks remain in the disk group for data integrity, and sufficient spare capacity is available, the storage system automatically uses the spare capacity to reconstruct the disk group. Disk group reconstruction does not require I/O to be stopped, so volumes can continue to be used while reconstruction is in progress.

If sufficient spare capacity is not available, reconstruction does not start automatically. For RAID levels other than ADAPT, to start reconstruction manually, replace each failed disk with a compatible disk. If the dynamic spares feature is not enabled, designate each replacement disk as a spare. If the dynamic spares feature is enabled, the storage system rescans the bus, finds the new disk, automatically designates it a spare, and starts reconstructing the disk group (as described in [Spares](#)).

For virtual storage, reconstruction of all disk groups uses a quick-rebuild feature. For more information, see [Quick rebuild](#).

For descriptions of LED states, such as for disk failure and reconstruction, see the related hardware documentation.

 **NOTE:** Reconstruction can take hours or days to complete, depending on the disk group RAID level and size, disk speed, host I/O activity, and other processes running on the storage system.

At any time after disk failure, you can remove the failed disk and replace it with a new disk of the same type in the same slot.

ADAPT reconstruction

Reconstruction of an ADAPT disk group is similar to reconstruction of a RAID-6 disk group, and can be impacted by host I/O activity and other processes running on the storage system.

ADAPT reconstruction differs from reconstruction of a RAID-6 disk group as follows:

- When one disk is failed, not all stripes will be degraded: there will be a mix of fault tolerant and degraded stripes.
- When two disks are failed, not all stripes will be critical: there will be a mix of fault tolerant, degraded, and critical stripes.
- Reconstruction will generally complete more quickly than for RAID-6.
- Reconstruction will start immediately without waiting for replacement of the failed disk.

NOTE: If a disk fails in an ADAPT disk group and is replaced by a new disk in the same slot as the failed disk, the disk group automatically incorporates the replacement disk into the disk group.

- Reconstruction will start on spare capacity already available in the ADAPT disk group.
- When there are critical stripes (and enough spare space), there will be two separate reconstruction phases: a first phase to repair critical stripes (to degraded state) and a second phase to repair the degraded stripes. Each phase will have its own start and end events. Because of the two-phase rebuild, ADAPT might take longer to reconstruct to fault-tolerant state than a critical RAID-6 running two-disk reconstruct. However, the first phase reconstruction of ADAPT, from critical state to degraded state, will be much faster. You can monitor reconstruction and rebalancing progress from the Activity panel.

If the ADAPT disk group has no spare space, the REFT (rebalance fault tolerant stripes) utility will run. As spare space is completely used, some stripes are critical, some are fault tolerant, and most are degraded. This utility attempts to rebalance stripe health away from the critical state and towards the degraded state. Stripes that are fault tolerant give up one of their disks, making them degraded. This disk capacity is then used to make a critical stripe zone degraded. It is recommended that spare space is added to the pool by either replacing failed disks or expanding the ADAPT disk group, and never to let spare space run out. However, if spare space is lost, the REFT utility attempts to give the ADAPT disk group the best redundancy across the whole disk group.

NOTE: Rebalancing-applicable only to ADAPT-will commence on the newly replaced disk. Use cases for rebalancing are described below:

- If the failed disk is replaced immediately, such that all stripe zones are fault tolerant, then only rebalancing occurs.
- If the failed disk is replaced later, and more disks have failed (such that there is limited or no spare space), then multiple stripe zones have likely become degraded or critical. Reconstruction will be followed by rebalancing.
- If no default spare space was selected, then reconstruction will occur without subsequent rebalancing

Quick rebuild

Quick rebuild is a method for reconstructing virtual disk groups that reduces the time that user data is less than fully fault-tolerant after a disk failure in a disk group. Taking advantage of virtual storage knowledge of where user data is written, quick rebuild first rebuilds the data stripes that contain user data.

Typically, storage is only partially allocated to volumes so the quick-rebuild process completes significantly faster than a standard rebuild. Data stripes that have not been allocated to user data are reconstructed in the background, using a lightweight process that allows future data allocations to be more efficient

Quick rebuild applies to all RAID levels, including ADAPT. For an ADAPT disk group, depending on how much space is actually allocated, quick rebuild can be faster than RAID rebuild

Within a few minutes after a quick rebuild completes, a scrub starts on the disk group. For more information about disk-group reconstruction, see [Reconstruction](#).

Updating firmware

Controller modules, expansion modules, and disk modules contain firmware. Users must have a Manage role to update the disk or system firmware. The PowerVault Manager provides options for you to update system firmware, disk firmware, and to use an update server (**Maintenance > Firmware**).

Related information

[Updating system firmware](#)

[Updating disk firmware](#)

[Firmware panel](#)

[Best practices for updating firmware](#)

Managed logs

As the storage system operates, it records diagnostic data in several types of log files. The size of any log file is limited, so over time and during periods of high activity, these logs can fill up and begin overwriting their oldest data. Enabling the managed logs feature (**Settings > System > Properties > Managed Logs Properties**) allows log data to be transferred to a log-collection system, and store it for later retrieval before any data is lost. The log-collection system is a host computer that is designated to

receive the log data transferred from the storage system. The transfer does not remove any data from the logs in the storage system. This feature is disabled by default.

The managed logs feature can be configured to operate in *push mode* or *pull mode*:

- In push mode, when log data has accumulated to a significant size, the storage system sends notifications with attached log files via email to the log-collection system. The notification will specify the storage-system name, location, contact, and IP address, and will contain a single log segment in a compressed zip file. The log segment will be uniquely named to indicate the log-file type, controller (A or B), and the date/time of creation. The file name format is `logtype_yyyy_mm_dd__hh_mm_ss.zip`. To activate push mode, select the **Include Logs** check box in the **Settings > System > Properties > Managed Logs Properties** panel.
- In pull mode, when log data has accumulated to a significant size, the system sends notifications via email or SNMP to the log-collection system, which can then use FTP or SFTP to transfer the appropriate logs from the storage system. The notification will specify the storage-system name, location, contact, and IP address and the log-file type (region) that needs to be transferred. To activate pull mode, clear the **Include Logs** check box in the **Settings > System > Properties > Managed Logs Properties** panel.


The managed logs feature monitors the following controller-specific log files:

- Expander Controller (EC) log, which includes EC debug data, EC revisions, and PHY statistics
- Storage Controller (SC) debug log and controller event log
- SC crash logs, which include the SC boot log
- Management Controller (MC) log

Each log-file type also contains system-configuration information. The capacity status of each log file is maintained, as well as the status of what data has already been transferred. Three capacity-status levels are defined for each log file:

- **Need to transfer**—The log file has filled to the threshold at which content needs to be transferred. This threshold varies for different log files. When this level is reached:
 - In push mode, informational event 400 and all untransferred data is sent to the log-collection system.
 - In pull mode, informational event 400 is sent to the log-collection system, which can then request the untransferred log data. The log-collection system can pull log files individually, by controller.
- **Warning**—The log file is nearly full of untransferred data. When this level is reached, warning event 401 is sent to the log-collection system.
- **Wrapped**—The log file has filled with untransferred data and has started to overwrite its oldest data. When this level is reached, informational event 402 is sent to the log-collection system.

Following the transfer of log data in push or pull mode, the log capacity status is reset to zero to indicate that there is no untransferred data.

 **NOTE:** In push mode, if one controller is offline, then its partner will send the logs it has acquired from the offline controller along with its own logs.


Saving log data to a file

To help service personnel diagnose a system problem, you might be asked to provide system log data.

Using the PowerVault Manager, you can save the following log data to a compressed zip file:

- Device status summary, which includes basic status and configuration data for the system
- The event log from each controller
- The debug log from each controller
- The boot log, which shows the startup sequence, from each controller
- Critical error dumps from each controller, if critical errors have occurred

Logs do not include user data.

 **NOTE:** The controllers share one memory buffer for gathering log data and for loading firmware. Do not perform more than one log saving operation at a time, or to perform a firmware update operation while performing a log saving operation. Also, do not attempt to gather performance metrics while performing a log saving operation.


Alternative methods for obtaining log data are to use the Collect Logs action (**Maintenance > Support**) or the `get logs` command in the FTP or SFTP interface. These methods will transfer the entire contents of a log file without changing its capacity-status level. Use of Collect Logs or `get logs` is expected as part of providing information for a technical support request.

For information about using the FTP or SFTP interface, see [Using FTP and SFTP](#).

DNS settings

You can set a domain hostname for each controller module to identify it for management purposes by configuring settings in the DNS panel (**Settings > Network > DNS**). The DNS server name supports IPv4 and IPv6 formats, and the system supports a maximum of three DNS servers per controller. Configuring the storage system to communicate with a DNS server within your network allows network changes, such as frequent IP address changes in a DHCP environment, to occur without interrupting notifications sent by the system to users.

The controller will advertise the domain hostname to DNS servers, and the DNS servers will in turn create and advertise a fully qualified domain name (FQDN) for the controller by appending the domain hostname to the DNS domain string that identifies the controller. The hostname must differ for each controller.

 **NOTE:** Rules for a valid domain name:

- The maximum domain name length is 63 characters.
- The domain name can contain alphanumeric characters and hyphens, but not periods.
- The domain name must not begin with a number, hyphen, or period; nor should it end with a hyphen.
- The domain name is not case sensitive.

After a reachable DNS server is configured on the system, you can configure an SMTP server using a name such as *mysmtpserver.example.com*. Further, you could configure search domain *example.com* and SMTP server *mysmtpserver* and reach the same destination.

You must use this feature to configure DNS parameters before you configure system parameters in any environments where DNS will be required to resolve server names.

If the controller is able to look up the domain name from a DNS server, the FQDN for each controller is also shown.

Peer connections

A peer connection enables bi-directional communication between a local system and a remote system to transfer data between the two systems. Creating a peer connection requires a name for the peer connection and either an IP address of a single available iSCSI host port on the remote system, or a WWN of a single available FC host port on the remote system. Only iSCSI and FC host ports are used for the peer connection. Establishing a peer connection for virtual storage is required for replication. SAS host ports do not support peer connections.

The peer connection is defined by the ports that connect the two peer systems, as well as the name of the peer connection. The local system uses the remote address to internally run the query `peer-connection` CLI command. The results of the query are used to configure the peer connection.

The prerequisites to create a peer connection are:

- Both systems must have iSCSI or FC host ports. Ports at both ends of the connection must use the same protocol.
- Both systems must be connected to the same fabric or network via a switch; direct connection between storage systems is not supported.
- All host port addresses in both systems must be unique, even for ports not in use.
- If iSCSI CHAP is configured for the peer connection, the authentication must be valid.

You can create a maximum of four peer connections per storage system. However, only one peer connection is allowed to a particular remote system. Attempting to create a second peer connection to the same system will fail. While creating the peer connection, the local system receives information about all host ports on the remote system as well as the port health on remote systems. It also links host ports of the selected host port type on the local system to those reachable on the remote system, so all ports of that type are available as part of the peer connection. Once created, the peer connection exists on both the local and remote systems.

Replications use a bi-directional communication path between the systems when exchanging information and transferring replicated data. Because the peer connection is bi-directional, replication sets can be created from both systems with replication occurring from either direction. Due to the relationship between peer connections and replication sets, creating a peer connection is part of the process of creating a replication set.

To create a peer connection, create a replication set **Provisioning > Volumes > Data Protection > Add Data Protection > Remote Replication**. Select one or more volumes to add to the replication set, then follow the on-screen directions to establish a peer connection between the primary and secondary systems. Directions include steps to enter the port address for the secondary system, the connection name, and the username and password of a user with a manage role on the remote system.

If a single host port loses connectivity, event 112 will be logged. Because a peer connection is likely to be associated with multiple host ports, the loss of a single host port may degrade performance but usually will not cause the peer connection to be inaccessible.

Related information

- [Add data protection](#)
- [Replication](#)
- [CHAP and replication](#)
- [Peer connection settings](#)

Replication

Replication for virtual storage provides a remote copy of a volume, volume group, or snapshot on a remote system by periodically updating the remote copy to contain a point-in-time consistent image of a source volume. After an initial image has been replicated, subsequent replications only send changed data to the remote system. All replications, including the initial one, only replicate data that has been written as opposed to replicating all of the data from the source. Replication can be used for disaster recovery, to preserve data, and to back up data to off-site locations. It can also be used to distribute data.

For more information about replication for virtual storage, see [CHAP and replication](#).

Replication prerequisites

To replicate a volume, you must first create a peer connection or use an existing one, and create a replication set. A peer connection establishes bi-directional communication between a local and remote system, both of which must have FC or iSCSI ports and a virtual pool. The system establishes a peer connection by connecting a host port on the local system with a user-specified host port on the remote system, then exchanging information and setting up a long-term communication path in-band. Because the communication path establishes a peer connection between the two systems, replications can occur in either direction.

To verify that a host port address is available before creating a peer connection in the PowerVault Manager, specify a peer system IP address and then choose **Query Peer Connection (Settings > Peer Connections)**. Alternatively, use the query `peer-connection` CLI command. This command provides information about the remote system, such as inter-connectivity between the two systems and pool configuration. For more information on this command, see the CLI Reference Guide.

In the PowerVault Manager, creating a peer connection or selecting an existing one, is part of creating a replication set. After you create or select a peer connection, you can continue to create a replication set. A replication set specifies one or more volumes, multiple volumes in a volume group, or snapshots on one system of the peer connection, known as the primary system in the context of replication, to replicate across the peer connection. When you create a replication set, corresponding volumes are automatically created on the other system of the peer connection, known as the secondary system, along with the infrastructure needed for replication. The infrastructure consists of internal snapshots used for replication operations. A replication set for a volume consumes two internal snapshots each for the primary volume and the secondary volume if the queue policy is set to Discard, or three each if the queue policy is set to Queue Latest. A replication set for a volume group consumes two internal volume groups if the queue policy is set to Discard, or three if the queue policy is set to Queue Latest. Each internal volume group contains a number of volumes equal to the number of volumes in the base volume group.

Using a volume group for a replication set enables you to make sure that multiple volumes are synchronized at the same time. When a volume group is replicated, snapshots of all of the volumes are created simultaneously. In doing so, the volume group functions as a consistency group, ensuring consistent copies of a group of volumes. The snapshots are then replicated as a group. Even though the snapshots may differ in size, replication is not complete until all of the snapshots are replicated.

For a replication set, the term *primary* refers to the source volume and the system in which it resides, and the term *secondary* is used for the remote copy and the system in which it resides. The secondary volume is meant to be an exact copy of the primary volume from the last time that replication occurred. To guarantee that the contents from that point in time match, the secondary volume cannot be mapped, rolled back, or modified except through replication.

While you cannot modify the secondary volume, you can create a snapshot of the secondary volume that you can map, mount, roll back, and otherwise treat like any volume or snapshot. You can regularly take snapshots to maintain a history of the replications for backup or archiving, or enable snapshot history for the replication set. These snapshots also can be used in disaster recovery.

NOTE: Both systems in a peer relationship should run the same firmware version. If you want to create a peer connection between a system running newer firmware and a system running older firmware, log in to the newer system and run commands to create and modify peers from that system.

Related information

- [Peer connections](#)
- [Peer connection settings](#)
- [Deleting a peer connection](#)

Replication process

After you create a peer connection and replication set, you can then replicate volumes between the systems. The initial replication differs slightly from all subsequent replications in that it copies all of the allocated pages of the primary volume to the secondary volume. Depending on how large your source volume is and the speed of the network connection, this initial replication may take some time.

Each subsequent replication is completed by resetting one of the internal snapshots of the primary volumes to contain the contents last replicated and then resetting the other internal snapshot to the current primary volume contents and comparing the changes. The system writes any changes it finds on the internal primary snapshot to the internal secondary snapshot, after which the secondary volume is updated to contain the contents of the secondary snapshot. (This internal process happens automatically, and is not accessible to user control.)

The progress and status of the initial and subsequent replications are tracked and displayed. The timestamps for replication reflect the time zones of the respective systems. When viewed on a secondary system in a different time zone, for example, replication information will reflect the time zone of the secondary system.

You can initiate a replication manually or by using a schedule. When creating a schedule for a replication set, you cannot specify for replication to occur more frequently than once an hour.

Related information

- [Aborting a replication set](#)
- [Initiating or scheduling a replication set](#)
- [Suspending a replication](#)

Initial replication

The following illustration shows the internal processes that take place during the initial replication of a single volume.




Figure 1. Process for initial replication

A	User view	c	Peer connection
B	Internal view	d	Primary volume
a	Primary system	e	Secondary volume
b	Secondary system	-	-

The two internal snapshots for each volume on the primary and secondary systems all have distinct roles. For both systems, they are labeled S1 (Snapshot 1) and S2 (Snapshot 2) in the two figures above and below. When a replication set is created, the primary volume and its internal snapshots all contain the same data. The secondary volume and its internal snapshots do not contain any data. Between the time that the replication set was created and the initial replication occurs, it is possible that hosts have written additional data to the primary volume.

During initial replication, the following sequence takes place.

1. The user initiates replication on the primary system.
2. The snapshot, S1, of the primary volume contents, which might be different than when the replication set was created, is reset to the current contents of the volume.
3. The S1 data, which matches that of the primary volume, is replicated in its entirety to its S1 counterpart on the secondary system and replaces the data that the secondary system S1 contains.
4. The S1 contents on the secondary system replace the contents of the secondary volume. The contents of the primary and secondary volumes are now synchronized.

Subsequent replications

The following figure illustrates the internal process that take place in replications subsequent to the initial replication of a single volume.

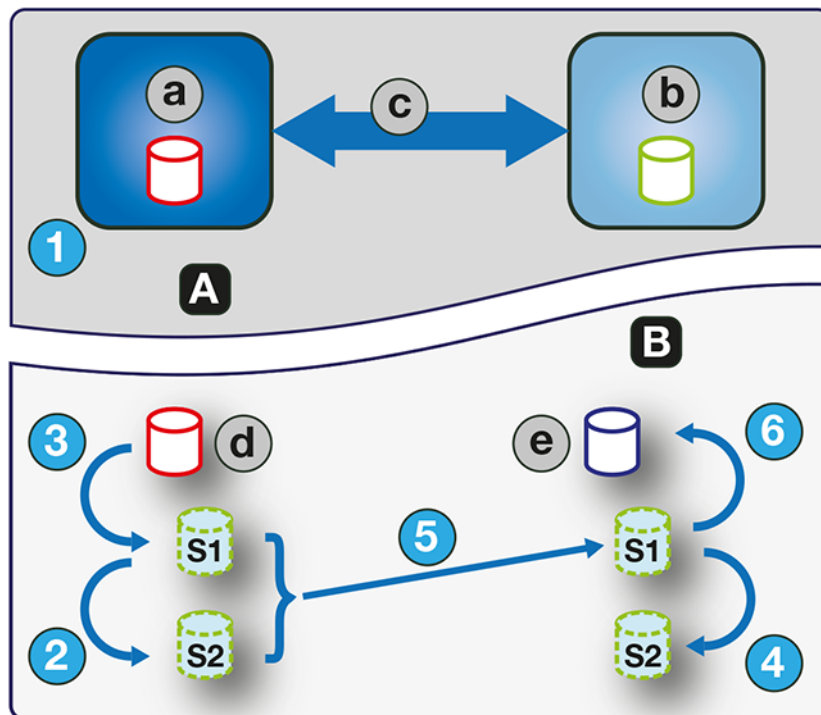


Figure 2. Process for subsequent replications

A	User view	c	Peer connection
B	Internal view	d	Primary volume
a	Primary system	e	Secondary volume
b	Secondary system	-	-

During the initial replication, the entire contents of the primary volume are replicated to the secondary volume. In subsequent replications, only data that is new or modified since the last replication operation is replicated. This is accomplished by comparing a snapshot of the primary volume data from the last replication with a current snapshot of the primary volume. With the exception of this comparison, the process for both the initial and subsequent replications is similar.

During replications subsequent to the initial replication, the following sequence takes place.

1. The user initiates replication on the primary system.
2. On the primary system, the S2 is reset to the current contents of S1. (The S2 contents can then be used for comparison during step 5.)
3. S1 is reset to the current primary volume contents on the primary system.
4. On the secondary system, S2 is reset to the current contents of S1.
5. The S1 contents on the primary system, which match that of the primary volume at the time the replication was initiated, are compared to the S2 contents on the primary system. Only the data that is the delta between S1 and S2 is replicated to its S1 counterpart on the secondary system, which is updated with the delta data. The data comparison and replication occur together.
6. The S1 contents on the secondary system replace the contents of the secondary volume. The contents of the primary and secondary volumes are now synchronized.

Internal snapshots

When first created from the primary volume, the internal snapshots consume very little space but will grow as data is written to the volume. Just as with any virtual snapshot, the amount of disk space used by an internal snapshot depends on the difference in the number of shared and unique pages between itself and the volume. The snapshot will not exceed the amount of disk space used by the primary volume. At most, the two internal snapshots together for each volume may consume twice the amount of disk space as the primary volume from which they are snapped.

Even though the internal snapshots are hidden from the user, they do consume snapshot space (and thus pool space) from the virtual pool. If the volume is the base volume for a snapshot tree, the count of maximum snapshots in the snapshot tree may include the internal snapshots for it even though they are not listed. Internal snapshots and internal volume groups count against system limits, but do not display.

Creating a virtual pool for replication

When you create a virtual pool, configure the size to allow for four times the anticipated size of the primary volume. This size will account for the primary volume, plus the same amount of space for each of the two internal snapshots, and possible queued replication. This is the maximum amount of space that you will need for replication. Also, for a pool on the primary system, allow additional space for other uses of the pool.

Setting up snapshot space management in the context of replication

Snapshot space management is the ability to control the number of snapshots and the amount of storage space they can consume in a pool.

When you create virtual volumes through the `create volume` and `create volume-set` CLI commands, you can set the retention priority for snapshots of the volume. If automatic deletion of snapshots is enabled, the system uses the retention priority of snapshots to determine which, if any, snapshots to delete. Snapshots are considered to be eligible for deletion if they have any retention priority other than `never-delete`. Snapshots are configured to be eligible for deletion by priority and age. The oldest, lowest priority snapshots are deleted first. Internal replication snapshots and snapshots that are mapped or are not leaves of a volume snapshot tree are ineligible for deletion. For more information on the `create volume` and `create volume-set` CLI commands, see the CLI Reference Guide

There are several factors to consider when managing snapshot space for the primary and secondary systems, especially when setting up the snapshot space and policies for the pool:

- Make sure that there is enough space in the pool for the base volume and replication-set needs. See guidelines provided in [Creating a virtual pool for replication](#).
- To adjust the snapshot space of the pool, increase the value of the limit parameter of the `set snapshot-space` CLI command. For more information on the `set snapshot-space` CLI command, see the CLI Reference Guide.

- To create more snapshot space, increase the size of the pool, add disk groups to the pool, or expand existing ADAPT groups. It may be necessary to insert additional disks if the system doesn't have any available.

If the snapshots are taking up more space than anticipated, consider the following approaches to address the issue:

- Limit the number of snapshots created by the replication-set.
- Fine-tune the overall snapshot space allocation and policies for the pool.

Limiting the number of snapshots created is the most straightforward approach. You can set the number of snapshots taken by the snapshot history feature via the CLI. Use the `set replication-set` CLI command, specifying the `snapshot-count` parameter. Reducing the snapshot-count number will save space. For more information on the `set replication-set` CLI command, see the CLI Reference Guide.

Fine-tuning the overall snapshot space allocation and policies for the pool is the other approach. You can adjust the snapshot space `limit` and `limit-policy` parameters of the pool to reduce the number of retained snapshots, or you can relax the `snapshot-retention-priority` parameter used with the `set replication-set` CLI command. Relaxing the `snapshot-retention-priority` parameter allows the pool to be more aggressive in deleting older, less important snapshots when there is pressure on the pool to free-up snapshot space. For more information on the `set replication-set` CLI command, see the CLI Reference Guide.

To monitor the snapshot space for virtual pools, use the `show snapshot-space` CLI command. To monitor the size of the internal snapshots, use the `show snapshots` CLI command with its `type` parameter set to `replication`. For more information on the `show snapshots` CLI command, see the CLI Reference Guide.

Replication and empty allocated pages

Deleting data from a volume can result in deallocation of pages on that volume. These pages consume space in the volume and in the pool. Pages deallocated before the initial replication will not be copied to the secondary volume. Pages deallocated since the last replication cause a page consisting of zeroes to be written to the secondary volume during replication.

This can result in a difference in the number of allocated pages between the primary and secondary volumes. A virtual storage background task automatically reclaims pages consisting of all zeroes, eventually freeing up the secondary volume snapshot space that these reclaimed pages consumed. Freeing up this space is not immediate and happens over a period of time.

Disaster recovery

The replication feature supports manual disaster recovery only. It is not integrated with third-party disaster recovery software. Since replication sets of virtual volumes cannot reverse the direction of the replication, carefully consider how the replicated data will be accessed at the secondary backup site when a disaster occurs.

NOTE: Using a volume group in a replication set ensures consistent simultaneous copies of the volumes in the volume group. This means that the state of all replicated volumes can be known when a disaster occurs since the volumes are synchronized to the same point in time.

Accessing the data while keeping the replication set intact

To continue replicating changed data from the primary data center system, you will need to keep the replication set intact.

While the data center system is down, you can access the data at the secondary backup system by creating a snapshot of the secondary volume or using the snapshot history snapshot. The snapshot can be mapped either read-only or read-write (but you cannot replicate the changes written to it back to the data center system using the existing replication set).

NOTE: If a system goes down but recovers, the data, peer connection, and replication sets should be intact and replication can resume normally.

To temporarily access data at the backup site

Follow these steps to access data at the backup site temporarily.

Steps

1. Create a snapshot of the secondary volume or use a snapshot history snapshot.
2. Map the snapshot to hosts.


3. When the data center system has recovered, delete the snapshot.

Accessing the data from the backup system as if it were the primary system

If you do not think the data center system can be recovered in time or at all, then you will want to temporarily access the data from the backup system as if it were the primary system.

You can again create a snapshot of the secondary volume and map that to hosts, or delete the replication set to allow mapping the secondary volume directly to hosts. Deleting the replication set means the secondary volume becomes a base volume and is no longer the target of a replication. Should the primary volume become available and you want to use it as is in preparation for another disaster, a new replication set with a new secondary volume must be created. Deleting the replication set also enables cleaning up any leftover artifacts of the replication set.

In an emergency situation where no connection is available to the peer system and you do not expect to be able to reconnect the primary and secondary systems, use the local-only parameter of the `delete replication-set` and `delete peer-connection` CLI commands on both systems to delete the replication set and peer connection. Do not use this parameter in normal operating conditions. For more information, see the CLI documentation. Other methods for deleting replication sets and peer connections will most likely be ineffective in this situation.

 **NOTE:** While deleting the peer connection for the replication set is unnecessary for making the secondary volume mappable, if you think that it will no longer be operable in the future, delete it when deleting the replication set.

Disaster recovery procedures

In a disaster recovery situation, you might typically perform the tasks in the following order:

Steps

1. Transfer operations from the data center system to the backup system (failover).
2. Restore operations to the data center system when it becomes available (failback).
3. Prepare the secondary system for disaster recovery.

Manually transfer operations from the data center system to the backup system

Steps

1. Create a snapshot of the secondary volume, use a snapshot history snapshot, or delete the replication set.
2. Map the snapshot or the secondary volume, depending on the option that you choose in step 1, to hosts.

Restore operations to the data center system

Steps

1. If the old primary volume still exists on the data center system, delete it. The volume cannot be used as the target—a new “secondary” volume will be created and deleting it will free up available space.
2. Create a peer connection between the backup system and the data center system, if necessary.
3. Create a replication set using the volume or snapshot on the backup system as the primary volume and the data center system as the secondary system.
4. Replicate the volume from the backup system to the data center system.

Prepare the backup system for disaster recovery after the replication is complete

Steps

1. Delete the replication set.
2. Delete the volume on the backup system. The volume cannot be used as the target of a replication and deleting it will free up space.
3. Create a replication set using the volume on the data center system as the primary volume and the backup system as the secondary system.

4. Replicate the volume from the data center system to the backup system.

Creating a replication set

A replication set specifies the components of a replication. You create replication sets using the Data Protection Configuration wizard.

You can access this wizard from the **Provisioning > Volumes** panel. For more information, see [Add data protection](#).

Performing this action creates the replication set and the infrastructure for the replication set. For a selected volume, snapshot, or volume group, the action creates a secondary volume or volume group and the internal snapshots required to support replications. By default, the secondary volume or volume group and infrastructure are created in the pool corresponding to the one for the primary volume or volume group (A or B). Optionally, you can select the other pool.

If a peer connection is not already created, the Data Protection Configuration wizard prompts you to create one. A replication set can specify only one peer connection and pool. When creating a replication set, communication between the peer connection systems must be operational during the entire process.

If a replication set is deleted, the internal snapshots created by the system for replication are also deleted. After the replication set is deleted, the primary and secondary volumes can be used like any other base volumes or volume groups.

Primary volumes and volume groups

The volume, volume group, or snapshot that will be replicated is called the primary volume or volume group. It can belong to only one replication set.

Using a volume group for a replication set enables you to make sure that the contents of multiple volumes are synchronized at the same time. When a volume group is replicated, snapshots of all of the volumes are created simultaneously. In doing so, it functions as a consistency group, ensuring consistent copies of a group of volumes. The snapshots are then replicated as a group. Though the snapshots may differ in size, replication of the volume group is not complete until all of the snapshots are replicated.

Secondary volumes and volume groups

When the replication set is created—either through the CLI or the PowerVault Manager—secondary volumes and volume groups are created automatically. Secondary volumes and volume groups cannot be mapped, moved, expanded, deleted, or participate in a rollback operation. If you need access to the data on secondary volumes, create a snapshot of the secondary volume or volume group and use the snapshot for mapping and accessing data.

Queuing replications

You can specify the action to take when a replication is running and a new replication is requested.

- **Discard**—Discard the new replication request.
- **Queue Latest**—Take a snapshot of the primary volume and queue the new replication request. If the queue contained an older replication request, discard that older request. A maximum of one replication can be queued. This is the default.

If the queue policy is set to **Queue Latest** and a replication is running and another is queued, you cannot change the queue policy to discard. You must manually remove the queued replication before you can change the policy using the `clear replication-queue` CLI command. For more information about this command, see the CLI Reference Guide.

Maintaining replication snapshot history

A replication set can be configured to maintain a replication snapshot history. As part of handling a replication, the replication set will automatically take a snapshot of the primary and/or secondary volume, thereby creating a history of data that has been replicated over time. This feature can be enabled for a secondary volume or for a primary volume and its secondary volume, but not for a volume group.

When this feature is enabled:

- For a primary volume, when a replication starts it will create a snapshot of the data image being replicated.

- For a secondary volume, when a replication successfully completes it will create a snapshot of the data image just transferred to the secondary volume. (This is in contrast to the primary volume snapshot, which is created before the sync.) If replication does not complete, a snapshot will not be created.
- You can set the number of snapshots to retain from 1 through 16, referred to as the snapshot retention count. This setting applies to management of snapshots for both the primary and secondary volume and can be changed at any time. Its value must be greater than the number of existing snapshots in the replication set, regardless of whether snapshot history is enabled. If you select a snapshot retention count value that is less than the current number of snapshots, an error message appears. Thus, you must manually delete the excess snapshots before reducing the snapshot count setting. When the snapshot count is exceeded, the oldest unmapped snapshot will be discarded automatically.
- You set the basename of the snapshot. The snapshots are named `basename_SNnnnn` where SN denotes snapshot and `_nnnn` starts at 0001 and increments for each subsequent snapshot. If primary volume snapshots are enabled, snapshots with the same name will exist on the primary and secondary systems. The snapshot number is incremented each time a replication is requested, whether or not the replication completes — for example, if the replication was queued and subsequently removed from the queue.
- If the replication set is deleted, any existing snapshots automatically created by snapshot history rules will not be deleted. You will be able to manage those snapshots like any other snapshots.
- If you begin keeping a snapshot history after a number of replications have occurred, the snapshot number in the name will reflect the total number of replications that have occurred.
- Manually creating a snapshot will not increase the snapshot count associated with the snapshot history. Manually created snapshots are not managed by the snapshot history feature. The snapshot history feature generates a new name for the snapshot that it intends to create. If a volume of that name already exists, the snapshot history feature will not overwrite that volume. Snapshot numbering will continue to increment, so the next time the snapshot history feature runs, the new snapshot name will not conflict with that existing volume name.
- The snapshot basename and snapshot retention count settings only take effect when snapshot history is set to secondary or both, although these settings can be changed at any time.
- A mapped snapshot history snapshot will not be deleted until after it is unmapped.
- A snapshot created by this feature is counted against the system-wide maximum snapshots limit, with the following result:
 - If the snapshot count is reached before the system limit then the snapshot history is unchanged.
 - If the system limit is reached before the snapshot count then the snapshot history stops adding or updating snapshots.
- The snapshot space management feature, accessible only through the CLI, enables users to monitor and control the amount of space that snapshots can consume in a pool. In addition to configuring a snapshot space limit, you can also specify a limit policy to enact when the snapshot space reaches the configured limit. An event is logged, and if the policy is set to delete, automatic deletion of snapshots occurs. If automatic deletion is triggered, snapshots are deleted according to their configured retention priority.
- You can set the retention priority for snapshots to the following. In a snapshot tree, only leaf snapshots can be deleted automatically.
 - **never-delete.** Snapshots will never be deleted automatically to make space. The oldest snapshot in the snapshot history will be deleted once the snapshot count has been exceeded. This is the default.
 - **high.** If snapshot space is exhausted, retained snapshots can be deleted after all eligible (system-wide) medium- priority snapshots have been deleted.
 - **medium.** If snapshot space is exhausted, retained snapshots can be deleted after all eligible (system-wide) low-priority snapshots have been deleted, regardless of retention schedule.
 - **low.** If snapshot space is exhausted, retained snapshots can be deleted at any time, regardless of retention schedule.

When this option is disabled, snapshot history will not be kept. If this option is disabled after a replication set has been established, any existing snapshots will be kept, but not updated.

CHAP and replication

If you want to use Challenge Handshake Authentication Protocol (CHAP) for the iSCSI connection between peer systems, see the procedure below to set up CHAP. Make sure that you configure both systems in this way. In a peer connection, both systems will alternately act as an originator (initiator) and recipient (target) of a login request.

If only one system has CHAP enabled and the two systems do not have CHAP records for each other, or the CHAP records have different secrets, the system with CHAP enabled will be able to modify the peer connection. However, it will be unable to perform any other replication operations, such as creating replication sets, initiating replications, replicating snapshots, or suspending replication operations. The system that does not have CHAP enabled will be unable to perform any replication operations, including modifying and deleting the peer connection. For full replication functionality for both systems, set up CHAP for a peer connection.

If the two systems have CHAP records for each other with the same secret, they can perform all replication operations whether or not CHAP is enabled on either system. In other words, even if CHAP is enabled on neither system, only one system, or both systems, either system can work with peer connections, replication sets, and replications.

If you want to use CHAP for the iSCSI connection between peer systems, see the following procedure to set up CHAP. In a peer connection, both systems will alternately act as an initiator and target of a login request. Peer connections support one-way CHAP only.

Set up CHAP for a peer connection using the CLI

To set up CHAP for a peer connection (using the CLI):

Steps

1. If you have not already configured CHAP, run `query peer-connection` from either the local system or the remote system to ensure that they have connectivity.
2. If you have an existing peer connection, stop I/O to it.
3. On the local system, use the `create chap-record` command to create a CHAP record for one-way CHAP to allow access by the remote system.
4. On the remote system, use the `create chap-record` command to create a CHAP record for one-way CHAP to the local system. Note that the same CHAP record used from the local system may also be used here but the configuration is still one-way CHAP.
5. On each system, enable CHAP by running: `set iscsi-parameters chap on`


 **CAUTION: Enabling or disabling CHAP will cause all iSCSI host ports in the system to be reset and restarted. This may prevent iSCSI hosts from being able to reconnect if their CHAP settings are incorrect.**

6. Wait one minute for the commands to complete before attempting to use the peer connection.
7. Run `query peer-connection` from the local system and then from the remote system to ensure communication can be initiated from either system.
 - If both succeed, you can create, set, or perform replication on that peer connection.
 - If either fails, it is likely that you must fix a CHAP configuration issue and then repeat these steps as appropriate. If you need to modify a CHAP record, use the `set chap-record` command.

Full disk encryption

Full Disk Encryption (FDE) is a method by which you can secure the data residing on the disks. It uses self-encrypting drives (SED), which are also referred to as FDE-capable disks. When secured and removed from a secured system, FDE-capable disks cannot be read by other systems.

The ability to secure a disk and system relies on passphrases and lock keys. A passphrase is a user-created password that allows users to manage lock keys. You can enable FDE protection by setting the FDE passphrase the system uses to write to and read from FDE-capable disks (**Settings > System > Security**). From the passphrase, the system generates the lock key ID that is used to secure the FDE-capable disks. If the system is unable to interpret the lock key on the FDE-capable disk, the encrypted data on the disk is inaccessible.

 **NOTE:** Be sure to record the passphrase as it cannot be recovered if lost.

A lock key is generated by the system, based upon the passphrase, and manages the encryption and decryption of data on the disks. A lock key is persisted on the storage system and is not available outside the storage system.

Data that was present on the system before it was secured is accessible in the same way it was when the system was unsecured. However, if a disk is transferred to an unsecured system or a system with a different passphrase, the data is not accessible.

Clearing the lock keys and power cycling the system denies access to data on the disks. Clear lock keys only when the system will not be under your physical control.

If the lock keys are cleared while the system is secured, the system will enter the FDE lock-ready state, in preparation for the system being powered down and transported. After the system has been transported and powered up, the system and disks will enter the secured, locked state; disks will be in the UNUSABLE state. Pools and disk-groups will be unavailable. All data on the disks is inaccessible until the system is secured with the original passphrase and lock key ID.

A system and the FDE-capable disks in the system are initially unsecured but can be secured at any point. Until the system is secured, FDE-capable disks function exactly like disks that do not support FDE.

FDE operates on a per-system basis, not a per-disk group basis. To use FDE, all disks in the system must be FDE-capable.

CAUTION: Do not change FDE configuration settings while running I/O. Temporary data unavailability may result, and the proper setting of lock keys from the passphrase could potentially be impacted.

Secured disks and systems can be repurposed. You can repurpose a system to erase all data on the system and return its FDE state to unsecured. You can repurpose a disk that is no longer part of a disk group. After a disk is repurposed in a secured system, the disk is secured using the system lock key ID and the new encryption key on the disk, making the disk usable to the system. Repurposing a disk in an unsecured system removes all associated lock keys and makes that disk available to any system.

CAUTION: Repurposing a disk changes the encryption key on the disk and effectively deletes all data on the disk. Repurpose a disk only if you no longer need the data on the disk.

NOTE: If you insert an FDE disk into a secured system and the disk does not come up in the expected state, perform a manual rescan. See [Rescanning disks](#).

Rescanning disks

A rescan (**Maintenance > Hardware > Actions**) forces a rediscovery of disks and enclosures in the storage system. If both storage controllers are online and can communicate with both expansion modules in each connected enclosure, a rescan also reassigns enclosure IDs to follow the enclosure cabling order of controller A. For further cabling information, refer to your product hardware documentation.

You might need to rescan disk channels after system power-up to display enclosures in the proper order. The rescan temporarily pauses all I/O processes, then resumes normal operation. It can take up to two minutes for enclosure IDs to be corrected.

You do not have to perform a manual rescan after inserting or removing non-FDE disks. The controllers automatically detect these changes. When disks are inserted, they are detected after a short delay, which allows the disks to spin up.

Clearing disk metadata

You can clear metadata from a leftover disk to make it available for use. This action is accessible via **Maintenance > Hardware > Disk > Actions** when a leftover (LEFTOVR) disk is selected. Choosing this action clears metadata only from leftover disks. If you specify disks that are not leftovers, the disks are not changed.

CAUTION: Consider the following points before clearing disk metadata:

- Only use this action when all disk groups are online and leftover disks exist. Improper use of this action may result in data loss.
- Do not use this action when a disk group is offline and one or more leftover disks exist.
- Do not use this action on disks that have gone leftover due to disk errors.
- If you are uncertain whether to use this action, contact technical support for assistance.

Each disk in a disk group has metadata that identifies the owning disk group, the other disks in the disk group, and the last time data was written to the virtual pool. The following situations cause a disk to become LEFTOVR.

- The disks timestamps do not match so the system designates members having an older timestamp as leftovers.
- A disk is not detected during a rescan, then is subsequently detected.
- A disk in a disk group is logically or physically removed from the system, and is later returned after the system has noted its removal.

When a disk becomes a leftover, the following changes occur:

- The disk health becomes Degraded and its usage value becomes LEFTOVR.
- The disk is automatically excluded from the disk group, causing the health of the disk group to become Degraded or Fault, depending on the RAID level.
- The fault LED on the disk is illuminated amber.

If a compatible spare is available, and the health of the disk group is Degraded or Critical, the disk group will use it to start reconstruction. When reconstruction is complete, you can clear the leftover metadata on the disk. Clearing the metadata deletes all data on the disk and change the disk health status to OK and its usage value to AVAIL. The disk may become available for use in a new disk group.

NOTE:

- If a spare is not available to begin reconstruction, or reconstruction has not completed, keep the leftover disk so that you will have an opportunity to recover its data
- ADAPT disk groups do not use spares for reconstruction.

NOTE: Disk health considerations:

- Before clearing metadata from a LEFTOVR disk to reuse it, check whether the disk previously reported excessive media errors. If so, the disk is probably not safe to use, and should be replaced.
- If the metadata on the disk has been cleared, verify that the disk health status is OK.
- When rebuilding a disk group, do not use an unhealthy disk from which metadata has been cleared

Data protection during failover to a single controller

The system can operate with a single controller if its partner has gone offline or has been removed but this configuration no longer provides redundancy.

When the ME5 system fails over to a single controller, the system switches to write-through caching mode. In write-through mode, the controller waits to notify the host that it has received a write until after the data is written to disk. This configuration guards against data loss during the failover.

If you set the caching policy to write-back mode, data to be written is held in the NVM cache. If the remaining controller fails while in write-back mode, unwritten cache data likely exists. The same is true if the controller or the enclosure of the target volume is powered off without a proper shutdown. Data remains in the controller cache and associated volumes will be missing that data on the disk.

- If the controller can be brought back online long enough to perform a proper shutdown and the disk group is online, the controller should be able to write its cache to disk without causing data loss.
- If the controller cannot be brought back online long enough to write its cache data to disk, contact technical support.

For more information about volume cache options, see [Volume cache options](#). To edit volume cache options, select **Provisioning > Volumes** and view the volume slide-over panel.

SupportAssist

SupportAssist provides an enhanced support experience for ME5 Series storage systems by sending configuration and diagnostic information to technical support.

Personal or business data is not collected in this process. SupportAssist only collects system state information including configuration, event notifications, and system diagnostic information. With SupportAssist, your information is safe during transport and storage. Data is transported securely with 256-bit encryption, secure web ports, and via a firewall with protected one-way communication from your site to Dell while your device usage and login credentials are never collected or leave your site.

If you need help with an issue and need to call technical support, they can access information about your storage system that was sent by SupportAssist. This feature enables technical support to assist you, without having to wait for configuration and diagnostic data to be collected and sent to technical support.

SupportAssist data

The configuration and diagnostic information that is sent by SupportAssist includes the following:

- System logs
- Events
- Topology
- Test connectivity
- CloudIQ configuration
- CloudIQ statistics

NOTE: The data that SupportAssist sends does not provide technical support with the information that is needed to connect to an ME5 Series storage system, because passwords are not transmitted.


About CloudIQ

CloudIQ provides storage monitoring and proactive service, giving you information tailored to your needs, access to near real-time analytics, and the ability to monitor storage systems from anywhere at any time. CloudIQ simplifies storage monitoring and service by providing:

- Proactive serviceability that informs you about issues before they impact your environment.
- Centralized monitoring across your environment, using a dashboard that aggregates key information such as system health scores, performance metrics, and current capacity and trends.

CloudIQ requires the following:

- SupportAssist must be enabled on the ME5 Series storage system.
- The ME5 Series storage system must be configured to send data to CloudIQ.






 **NOTE:** For more information about CloudIQ, contact technical support or visit the [CloudIQ](#) product page.


Event history

If you are having a problem with the system, review the event history (**Maintenance > Support > Event History**) to view event details and recommended actions before calling technical support. Information shown might enable you to resolve the problem.

All events are logged, regardless of notification settings. For information about notification settings, see [Notification settings](#). The event history table lists a collapsed view of the most recent events logged by either controller module, up to 1000. For each event, the table shows the date and time when the event occurred (with one-second granularity), the severity, which controller module logged the event, the event code, and a message. For information about using tables, see [Tips for using tables](#).

Table 9. Event severity icons and meanings

Icon	Severity	Meaning
	Critical	A failure occurred that may affect data integrity, system stability, or cause a controller to shut down. Correct the problem immediately.
	Error	A failure occurred that may affect data integrity or system stability. Correct the problem as soon as possible.
	Warning	A problem occurred that may affect system stability but not data integrity. Evaluate the problem and correct it if necessary.
	Informational	A configuration or state change occurred, or a problem occurred that the system corrected. No action is required
	Resolved	A condition that caused an event to be logged has been resolved. No action is required.

When reviewing the event history, look for recent Critical, Error, or Warning events. For each, click the  icon to view additional information and recommended actions. Follow the recommended actions to resolve the problems. To take action to correct an event, see the [Alerts panel](#).

Audit logs

User login, logout, and actions through all interfaces will be recorded in an audit log that is accessible from **Maintenance > Support > Audit Log Activity**. For actions that change storage system parameters, the audit log will contain the timestamp, username, and actions that were run as well as the status code returned by that action.

The audit log will include operations performed using the PowerVault Manager, CLI, and FTP/SFTP protocols, but will not contain specific value changes, such as old and new settings.

Audit logs record host IP address information for all interfaces, as well as SNMP SET requests. Each controller maintains its own audit log, and each audit log can contain up to 2MB of data. The system will maintain 2MB of audit log data per controller and the log will wrap once the 2MB limit has been reached.

Audit log data will persist after the restore defaults CLI command is run and is not mirrored to the partner controller. In a failover scenario, the failed controller audit log cannot be retrieved. Audit log data is cleared during factory refurbishment. When you download controller logs, audit logs will be included. Audit logs are not associated with the managed logs feature.

System metrics

There are two types of metrics: historical and dynamic.

The sampling interval for a dynamic metric is 5 seconds while that of a historical metric is 5 minutes. The system samples every metric at a dynamic sampling interval. The last few minutes of samples are kept in cache and cache is flushed automatically at historical sampling intervals with calculated maximum, minimum, and average data points retained in the historical database. Only a few dynamically sampled user-selected data series are retained. For such user selected data series, several hours of data points are available. The Performance panel provides charting and data comparison capabilities for metrics that are maintained in the system. For more information, see the [Performance panel](#).

Effect of failover on metrics

Controller A is the active metrics server by default. It has the one active metrics database on the storage system. Controller B does not have a metrics database, and simply forwards samples to controller A as they are received from Storage Controller B. This single-metrics database exists in RAM only. It is not redundantly maintained, and is lost if controller A goes offline.

If controller A goes offline, all current metrics samples are lost. Controller B becomes active, creates an empty metrics database on the memory card, and begins storing and persisting samples to cover for controller A. When controller A is back online, controller B clones its current database to controller A. Controller A then initializes its metrics database with the contents of those files.

Dashboard

Use the PowerVault Manager dashboard to monitor the system and see an overview of system status.

Each panel on the dashboard has a compact view, which is the summary information that appears on the dashboard itself, and an expanded view that provides more detailed information about the topic and lets you take specific actions based on the information.

Topics:

- [Alerts panel](#)
- [Capacity panel](#)
- [Performance panel](#)
- [Activity panel](#)

Alerts panel

Use the Alerts panel to monitor system health and performance issues and to track and acknowledge the resolution of these issues.

For alert severity level meanings, see [Areas of the interface](#).

Compact view

The compact view on the dashboard provides a snapshot of the overall system health, including a count of health alerts, information alerts, and unacknowledged alerts.

- Click **System Health | Health Alerts** to view a scrollable list of unresolved health alerts that are affecting system health.
- Click **Information Alerts** to view a scrollable list of unresolved informational alerts that notify you of actions to take to improve system performance.
- Click **Alerts to Acknowledge** to view a scrollable list of all resolved and unresolved health and information alerts that need acknowledgment. Alerts are removed from the list after they are acknowledged and resolved.
- Click **Acknowledge** or the slide-over arrow next to it to open the Alerts panel expanded view.


Expanded view

The expanded view shows the scrollable Active Alerts table and lists detailed alert information. The alerts that display are dynamic and based on the type of alerts you want to view (Health Alerts, Information Alerts, Alerts to Acknowledge, or History).

Click the name of an alert type to view more information. For each alert, the table shows:

- How long the alert has been active.
- The alert severity.
- The affected system component.
- A description of the problem.
- Whether the alert has been resolved (Alerts to Acknowledge table).
- Whether the alert has been acknowledged (Health Alerts, Information Alerts, and Alerts to Acknowledge tables).

The Alerts to Acknowledge view shows only alerts that have not been acknowledged.

Click  to see additional detail:

- The date and time when the issue was detected.
- The date and time when the issue was resolved, if applicable.
- A more detailed description of the problem, if available
- One or more recommended actions to resolve the problem, if applicable

- A View on System link to view component details, for certain types of events

To acknowledge an alert, click **Acknowledge**. If the alert is resolved, the entry is removed from the Alerts to Acknowledge table and will only appear in the History table. Unresolved alerts remain in the Health Alerts or Information Alerts list until they are resolved.

To acknowledge all alerts, click **Acknowledge All Alerts** in the Alerts to Acknowledge table. Acknowledging all alerts removes all current alerts from the list and place them in the History table. Unresolved alerts remain in the Health Alerts or Information Alerts list until they are resolved.

Click **History** to view a scrollable, read-only log of conditions that have caused the system to generate alerts. Use it to identify components that repeatedly cause alerts. This information can provide valuable troubleshooting information. Historical alerts remain in the History table until it reaches its threshold limit of 512, after which the table content wraps.

Capacity panel

Use the Capacity panel to monitor storage usage trends at the disk, disk group and volume level.


Compact view

The content displayed for this view depends upon the selected storage type.

Linear storage

The compact view on the dashboard shows object counts for volumes, disks, and disk groups in the top section. The bottom section provides a spacemeter with a capacity graph for disk usage and a capacity graph for volume usage. Each capacity graph uses a segmented horizontal bar.

- Disk usage is labeled with text and coded as Used (green), Unusable (amber), or Available (gray).
- Volume usage is labeled with text and coded as Allocated (green), Unallocated (gray), or Unavailable (two-tone gray diagonally slanted stripes).
- Spares are shown only if the number is non-zero.

 **NOTE:** The space available for volume allocation equals the size of the disk-group(s) minus overhead. Overhead includes spares, RAID overhead, and unusable or unused disks. Relative to the two capacity graphs, the text label for Disk Usage (Used) matches the text label for Volume Usage (Total Size).

Virtual storage


The compact view on the dashboard shows a capacity graph for each configured pool. Each capacity graph uses a segmented radial bar.

The radial bar segments are color-coded as follows:

- Green—Pool space that is allocated to volumes in the pool.
- Grey—Unallocated pool space that remains available to the system.

General

Text labels on each capacity graph provide clarification and report metrics for capacity use and availability.

Click  to see the expanded view.

Expanded view

The content displayed for this view depends upon the selected storage type.

Linear storage

For linear storage, the expanded view shows the following:

- Disk-group allocation (click the arrow to expand the row), with color-coding defined on-screen. Disk group allocation is represented as a segmented horizontal bar. Text labels reporting metrics differ per the given spacemeter:
 - ADAPT disk-group allocation spacemeter: allocated, spare capacity, overhead, and available.
 - Non-ADAPT disk-group allocation spacemeter: allocated, overhead, and available. A dedicated spares value reports the number of dedicated spares for the particular disk-group.
- Volumes for each disk group defined (click the arrow to expand the row). For each configured volume, the panel shows:
 - A filtering field with drop-down selection.
 - A table providing data for each volume: name, size, percentage of disk group, and attached hosts.


Virtual storage

For virtual storage, the expanded view shows the following:

- System level counts of volumes, snapshots, and disks that are in use versus system configuration limits.
- For each configured pool, the panel shows:
 - The capacity graph, with a link to **View I/O Workload**.
 - Counts of volumes and snapshots that are in use in that pool versus system configuration limits.
 - A pool allocation graph, with color-coding defined on-screen. Pool allocation is represented as a segmented horizontal bar. Text labels on the pool allocation graph report metrics for volume allocated, snapshot allocated, total committed pool allocation, and unused capacity.
 - A tiering capacity graph, with color-coding defined on-screen. Tiering capacity is represented as a segmented horizontal bar. Text labels on the tiering capacity graph report metrics for allocated and available capacity on each available tier.
- Volumes for each pool defined (click the arrow to expand the row). For each configured volume, the panel shows:
 - A filtering field with drop-down selection.
 - A table providing data for each volume: name, attached hosts, creation date, tree size, and capacity utilization.


Performance panel


Use the Performance panel to monitor system performance and statistics.

 **NOTE:** PowerVault Manager displays the Performance panel only when virtual pools are present. Linear pools are not supported in the Performance panel.

Compact view


The compact view on the dashboard shows consolidated metrics for up to eight graphs selected as favorites in the expanded view. By default, three graphs are shown: System IOPS (read and write), System Latency (read and write), and System Throughput (read and write).

 **NOTE:** Compact view does not display live metrics. If the Live option is selected in the expanded view, the graphs in the compact view display the most recent historical data for the last four hour period.


Click  to see the expanded view.

Expanded view

The expanded view shows more detailed information and provides options to view metrics within a select time range, add graphs, configure existing graphs, and to select up to eight graphs as favorites to show in the compact view. You can also choose to display live data up to the last 15 minutes.

 **NOTE:** Viewing live data is only available in the expanded view. It is not shown in compact view.

You can view and compare data-point values by hovering the cursor over the graph, which shows the values in a pop-up table.

To see average and peak values for data in a graph, click 

Performance metrics

The following system performance metrics and statistics are displayed on the Performance panel.

Table 10. Available performance metrics

Metric	Description	Storage Object ^{1,2}				
		Controller	Host port	Pool	System	Volume
IOPs	<ul style="list-style-type: none"> • Total—Sum of read IOPs and write IOPs. • Read—Number of I/Os per second. • Write—Number of I/Os per second 	✓	✓	CLI	✓	✓
Throughput	<ul style="list-style-type: none"> • Total—Sum of read bytes per second and write bytes per second. • Read—Number of bytes read per second. • Write—Number of bytes written per second 	✓	✓	CLI	✓	✓
Latency	<ul style="list-style-type: none"> • Average Response time—Average response time of an operation in microseconds. Operations include both reads and writes. • Max Response time—Sum of read maximum response time and write maximum response time. • Total—Sum of read and write average responses in microseconds. • Read Average—I/O read average response time in microseconds. • Write Average—I/O write average response time in microseconds. • Total—Sum of read and write maximum responses in microseconds. • Read Max—Maximum I/O read response time in microseconds. • Write Max—Maximum I/O write response time in microseconds. 	✓	✓	CLI	✓	✓
Average Queue Depth	<ul style="list-style-type: none"> • Read—Average number of pending read operations being serviced since the last sampling time. This value represents periods of activity only and excludes periods of inactivity. • Write—Average number of pending write operations being serviced since the last sampling time. This value represents periods of activity only and excludes periods of inactivity. 	-	✓	CLI	-	✓
Cache	<ul style="list-style-type: none"> • Read Ahead Ops—Number of times that read ahead pre-fetched data for host reads. 	✓	-	CLI	-	✓

Table 10. Available performance metrics (continued)

Metric	Description	Storage Object ^{1,2}				
		Controller	Host port	Pool	System	Volume
	<ul style="list-style-type: none"> • Small Destages—Number of partial stripe destages. (These tend to be very inefficient compared to full stripe writes.) • Write Cache Percent—Percentage of write cache currently being used in tenths of a percent. • Write Cache Space—Current size of write cache in 16KB chunks. 					

¹ Metrics for the Pool storage object are not graphed in the PowerVault Manager, but are accessible via the CLI.

² Metrics for the System storage object are synthesized from data captured by Controller storage objects in the storage system.

Collecting performance data

View performance metrics data for a specified period by choosing an option from the Timeframe drop-down list. Your selection determines your viewing window.


Live performance metrics show the current system performance within the last 15 minutes and are sampled every five seconds. The remaining options let you choose to display historical performance up to the last eight days and are sampled every five minutes.

The collection of metrics never stops, except in cases where the system resets. The list of metrics that are collected is defined for the system, and not an individual user. If one user redefines the metrics to track, all user sessions are affected.









Controller A retains all the metrics in memory regardless of whether you are logged in to the PowerVault Manager. If controller A is restarted, all metrics are lost. If controller B is restarted, metrics continue to be available.

Displaying performance data


Performance metrics are displayed in graphs, allowing for easy charting, data comparison, and analysis. Data points in each graph line up from top to bottom, allowing you to easily compare data points between graphs when hovering over them.

 **NOTE:** Clicking an object in the legend of a graph highlights the object data by hiding other objects.

The following options are available when in expanded view:

- Click the **Data Timeframe** dropdown to view metrics within a select time range. **Last 1 day** is the default.
 -  **NOTE:** Choosing Live refreshes the graphs every five seconds while in expanded view and displays up to the last 15 minutes of data in five second increments. Viewing live graphs can only be done in expanded view. Graphs shown in compact view reflect the date range that is selected in expanded view *except for the Live option*. When Live is selected, static graphs display the most recent historical data for the last four-hour period when you return to compact view.
- Click  to view average and peak data values.
- Click **Add Graph** to build a graph by selecting metrics from a predefined list. Graph settings are unique to each graph. The maximum number of objects that can be in any single graph is four (volumes, host ports, and so forth). Also, each object can be charted up to three times in the same graph for a comparison (for example, read vs write vs total operations).
- Click  to designate performance metric favorites (displayed in compact view). Default favorites include System IOPS, System Throughput, and System Latency (displaying both Read/Write).
- Click  or  to move the graph up or down in the display.
- Click  to access the configuration options for the graph.
- Click  to download historical performance metrics in CSV format.
- Click  to remove the graph from view. (The graph can be added back.)

The system retains historical data for one week, or since the last controller restart if that is more recent. By default, the graphs show the latest 100 data points. If a time range or count of samples returns more than 100 data points, adjacent data samples are combined so that only 100 data points are shown. These combined data points are displayed on the graph.



 **NOTE:** If a volume that was added to a Performance Metrics graph is deleted, the graph configuration must be manually edited to remove that volume. Otherwise, an error message instructing you to check and update graph settings appears in the graph display.

Activity panel

Use the Activity panel to monitor system activities that are currently running, those that are scheduled, and those that have been recently completed. Activities include Scheduled Tasks, Jobs, and Utilities.

Compact view

The compact view on the dashboard shows the most recent system activities that are in progress, are scheduled, or have completed.

- **In Progress**—Activities that have started but are not yet complete. This includes disk-group jobs such as initialization and scrub, and long-running application activities such as replicating volumes. A progress bar shows status using green fill and text reporting percentage complete.
 **NOTE:** The PowerVault Manager uses the last 1000 events to display the activity progress. If a long-running application activity start event is not found, it will report **Timestamp Unavailable**.
- **Scheduled**—Activities that users have scheduled to run. This includes enabling and disabling disk spin down (DSD); taking and resetting snapshots; copying volumes; and replicating volumes.
- **Recent**—System activities and user actions that have already completed. For example, a disk-group scrub job completed or a volume was deleted. Click  to see the expanded view.

Expanded view

The expanded view enables you to view In Progress, Scheduled, or Recent activities in a tabular format and act on them.

- **In Progress** table—shows all activities that are in progress, in reverse-chronological order by start time by default. For a selected activity, based on its type, you can view extra information or configuration associated with the activity; view or modify its schedule; or perform actions such as suspending, resuming, or aborting it.
- **Scheduled** table—shows all activities that are scheduled to run. For a selected activity, its schedule can be viewed or modified, and configuration specific to the scheduled operation can be modified.
- **Recent** table—shows a history of activity on the system, including both successful and unsuccessful operations. Where errors occur and more information can be gathered, you can select an entry and view details.

The **Filter By** drop-down list enables selection of activities to display and monitor.

For tips on using tables, see [Tips for using tables](#).

Provisioning

Use the Provisioning panel to view, create, and manage volumes and hosts, replication sets, and snapshots.

Topics:

- [Working with volumes](#)
- [Working with hosts](#)

Working with volumes

The Volumes panel (**Provisioning > Volumes**) provides options to create volumes, add data protection to volumes, and to display a list of existing volumes and snapshots on the system.


For information about creating volumes, see [Creating volumes](#).





If storage is not configured, the panel displays a prompt for you to add storage. Depending on the storage type selected:

- If virtual storage is configured, the top portion of the panel displays a color-coded graph for each pool that shows the amount of space on the system that is allocated to the base volume, allocated to snapshots, the total committed to volumes (reserved for the system but not written to), and unused space. Each named volume is shown in the Volumes table with columns for type, pool, size, attached hosts, capacity, and a slide-over for access to additional detail and actions.
- If linear storage is configured, each named volume is shown in the Volumes table with columns for type, disk group, size, attached hosts, and a slide-over for access to additional detail and actions

Volumes table

To perform an action on an existing volume or snapshot, select one or more volumes or snapshots and choose an option from the drop-down list.

Click the volume slide-over  to view volume or snapshot details and to perform more actions on a volume or snapshot.

- The Overview tab shows volume capacity usage, volume copy activity, and lists all of the volume settings. Click  to make changes to these settings. Click **Expand Volume** to expand the size of the volume. Click **Copy Volume** to copy the volume. Follow the on-screen directions for more details.
- The Snapshot tab displays the snapshots associated with the volume, along with any associated schedules. Click  to edit snapshot schedules. The Snapshot Tree table lists a hierarchical table of the snapshot tree for the selected volume. Click a volume or a snapshot of the volume, then choose an option from the drop-down list to perform an action. Click the parent volume or snapshot, then choose **Add Data Protection** and follow the on-screen directions to create local snapshots.
- The Attached Hosts tab displays a table listing the hosts attached to the selected volume or snapshot and lets you attach the volume or snapshot to a host or host group. Hover over icons in the list to see more information. Click  to edit permissions, LUN IDs, and ports. The table shows the following information about each attached host:
 - **Name**—The name of the attached host.
 - **Discovered**—Shows if the host is currently logged into the system.
 - **Redundant**—Shows if the host is logged into both controllers of the system.
 - **Mapped**—Shows if the volume is presented to both controllers.
 - **Permissions**—Displays the volume access permissions.
 - **LUN**—Displays the ID used to identify the volume on the host.
 - **Ports**—Displays the ports that the LUN is presented (the default is all ports).
 - **Unmap**—A clickable option allowing you to detach the host from the volume.
- The Replications tab displays the replication sets associated with the volume, along with any associated schedules. Click **Add Data Protection** to access the Data Protection wizard; click **Start Replication**, **Suspend Replication**, or **Remove Replication** to perform the action; click  to edit replication schedules and replication details.


Related information


[Deleting volumes and snapshots](#)
[Copying volumes or snapshots](#)
[Aborting a volume copy](#)
[Expanding volumes](#)
[Modifying volumes](#)
[Add data protection](#)
[Cache optimization mode](#)
[Optimizing read-ahead caching](#)
[Volume tier affinity](#)
[Attaching volumes to hosts](#)
[Detaching volumes from hosts](#)
[Rolling back virtual volumes](#)
[Creating snapshots](#)
[Resetting snapshots](#)
[Replication](#)
[Initiators, hosts, and host groups](#)






Data protection table


Click **Data Protection** in the Volumes table (**Provisioning > Volumes**) to display a list of existing volumes on the virtual storage system that are protected by a replication set or a snapshot. Click **Add Data Protection** to access the Data Protection wizard to protect your volumes using remote replication or local snapshots.

To perform an action on one or more volumes in the table, select the volumes and then choose an option from the drop-down list:

 **NOTE:** Some options require that only one volume is selected.

Click the slide-over  to perform more actions on a volume or snapshot.

- The Overview tab displays the volume settings. Click  to make changes to these settings. Click **Copy Volume** to copy the volume. Follow the on-screen directions for more details.
 **NOTE:** Snapshots cannot be expanded.
- The Snapshots tab displays the snapshots associated with the volume, along with any associated schedules. Click  to edit snapshot schedules. The Snapshot Tree table lists a hierarchical view of the snapshots for the selected volume. Click a snapshot, then choose **Add Data Protection** and follow the on-screen directions to create local snapshots, or choose an option from the drop-down list to perform an action:
 **NOTE:** When creating a snapshot, you will be prompted to set a snapshot schedule. This is the only place where you can schedule a snapshot using the PowerVault Manager. If you do not set a snapshot schedule when prompted, you will be unable to do so using the PowerVault Manager after the snapshot is created. For information about creating a snapshot schedule using the `set schedule` CLI command, see the CLI Reference Guide
- The Attached Hosts tab displays a table listing the hosts attached to the selected volume or snapshot and lets you attach the volume or snapshot to a host or host group. Hover over icons in the list to see more information. Click  to edit permissions, LUN IDs, and ports. The table shows the following information about each attached host:
 - **Name**—The name of the attached host.
 - **Discovered**—Shows if the host is currently logged into the system.
 - **Redundant**—Shows if the host is logged into both controllers of the system.
 - **Mapped**—Shows if the volume is presented to both controllers.
 - **Permissions**—Displays the volume access permissions.
 - **LUN**—Displays the ID used to identify the volume on the host.
 - **Ports**—Displays the ports that the LUN is presented (the default is all ports).
 - **Unmap**—A clickable option allowing you to detach the host from the volume.
- The Replications tab displays the replication set associated with the volume, along with any associated schedules. If there is no existing replication set for the volume, click **Add Data Protection** to access the Data Protection Configuration wizard;

if there is a replication set for the volume, click **Start Replication**, **Suspend Replication**, or **Remove Replication** to perform the action; click  to edit any replication schedules and replication details.

NOTE: The Status label shows the state of replication activity. For details about the possible replication states, see the on-screen tool tip next to the active state.

NOTE: Selecting **Start Replication** or **Remove Replication** will begin the action without any warning.

Related information

[Deleting volumes and snapshots](#)

[Copying volumes or snapshots](#)

[Add data protection](#)

[Rolling back virtual volumes](#)

[Creating snapshots](#)

[Resetting snapshots](#)

[Replication](#)

[Initiators, hosts, and host groups](#)

[Snapshots](#)

[Volumes and volume groups](#)

[Data protection with a single controller](#)

Creating volumes

You can add volumes to virtual pools and linear disk groups. Click **Create Volumes (Provisioning > Volumes)** to open the Create Volumes wizard to add volumes to a pool (virtual) or a volume to a disk group (linear).

Depending on the storage type:

- The top portion of the panel displays a color coded graph for each pool on the system or volume usage by disk group.
- The graph indicates:
 - Amount of space on the system that is allocated to volumes
 - Total volumes committed (reserved for the system but not written to) (virtual)
 - Space occupied by added volumes
 - Space required for the new volume you are creating
 - Unused available space (linear)
 - Space that is overcommitted if the Pool Overcommit setting is enabled (virtual)

Follow the on-screen directions to create one or more new volumes to add them to the table. Fields with a red asterisk are required. Choose **Continue** when you finish creating volumes. The wizard prompts you to attach the volume to a host or host group, or allows you to create the volumes and attach hosts or host groups later. Choose the former option to attach the volumes to new or existing hosts or host groups. Choose the latter option to create volumes that are not attached to hosts or host groups. New volumes are listed in the Volumes table.

You can create an individual volume or multiple volumes. For more information, see:

- [Volumes and volume groups](#)
- [Initiators, hosts, and host groups](#)


NOTE: Virtual volume sizes are aligned to 4.2-MB (4-MiB) boundaries. When a volume is created or expanded, if the resulting size would be less than 4.2-MB it will be increased to 4.2-MB; if the resulting size would be greater than 4.2-MB it will be decreased to the nearest 4.2-MB boundary. For the maximum volume size supported by the system, see System configuration limits.

NOTE: The PowerVault Manager supports stripe-aligned and non-stripe-aligned linear volumes. Rounding rules for protection levels are summarized:

- ADAPT protection level rounds to the nearest stripe zone size of 8-GiB for 8+2 and 16-GiB for 16+2 stripe widths.
- RAID 5/6/10/0 protection levels round to a stripe boundary. A stripe is defined as $(NumDataDrives \times ChunkSize)$:
 - NumDataDrives = DrivesInDiskGroup - NumParityDrives
 - NumParity drives is 1 for RAID 5 and 2 for RAID 6.
 - ChunkSize is defined by the user when the disk group is created.

Modifying volumes

Change the volume settings from the Volumes table (**Provisioning > Volumes**) by selecting the volume slide-over to access the Overview panel. Here you can expand volumes, copy volumes, modify the volume name, and select cache setting options and tier affinity settings. If the volume is not a snapshot or a secondary volume involved in replication, you can expand the volume size.

 **CAUTION:** Only change the volume cache settings if you fully understand how the host operating system, application, and host adapter move data so that you can adjust the settings accordingly.

Related information

[Volume cache options](#)

[Volume tier affinity](#)

Deleting volumes and snapshots

You can delete volumes and snapshots from the Volumes table:


- Volumes table (**Provisioning > Volumes**)
- Data Protection table (**Provisioning > Volumes > Data Protection**)
- Snapshots panel (**Provisioning > Volumes > slide-over > Snapshots**)


From the slide-over you can only delete the selected volume (the volume for which the slide-over is opened) and its children. Clicking the slide-over for the base volume enables deleting the entire tree.

 **CAUTION:** Deleting a volume or snapshot removes its host attachments and schedules and deletes its data.

Select a volume or snapshot, then choose **Delete Volumes** from the drop-down list. Follow the on-screen directions to complete the action. The following rules apply:

- You can select from 1 to 100 items (volumes, snapshots, or both) to delete.
- You cannot delete a volume that is part of a replication set.
- Ensure that hosts are not accessing the volumes or snapshots to be deleted.

 **NOTE:** To delete a volume with child snapshots, delete all of the child snapshots first.

 **NOTE:** To delete a volume that is part of a replication set, delete the replication set first.

Attaching volumes to hosts

Attach volumes to hosts from one of the following:


- Volumes table (**Provisioning > Volumes**)—Select the volume and choose **Attach to Hosts** from the drop-down list. Follow the on-screen directions to complete the action.
- Attached Hosts panel (**Provisioning > Volumes > slide-over > Attached Hosts**)—Follow the on-screen directions to complete the action.

 **NOTE:** From the slide-over you can only attach the selected volume (the volume for which the slide-over is opened).

Detaching volumes from hosts

You can detach volumes from hosts from the one of the following:


- Volumes table (**Provisioning > Volumes**)—Select a volume and choose **Detach from Hosts** from the drop-down list.
- Attached Hosts panel (**Provisioning > Volumes > slide-over > Attached Hosts > unmap**)
- Data Protection table (**Provisioning > Volumes > Data Protection > slide-over > Attached Hosts > unmap**)—Follow the on-screen directions to complete the action.

 **NOTE:** From the slide-over you can only detach the selected volume (the volume for which the slide-over is opened).

Expanding volumes


You can expand the size of a volume from one of the following:


- Volumes table (**Provisioning > Volumes**)—Select the volume and choose **Expand Volumes** from the drop-down list.
- Overview panel (**Provisioning > Volumes > slide-over > Overview > Expand Volume**)
- Data Protection table (**Provisioning > Volumes > Data Protection > slide-over > Overview > Expand Volume**)—Follow the on-screen directions to complete the action.

 **NOTE:** From the slide-over you can only expand the selected volume (the volume for which the slide-over is opened).

The top portion of the panel displays a color-coded graph for each pool on the system or volume usage by disk group. The graph indicates the amount of space on the system that is allocated to volumes; the total volumes committed (reserved for the system but not written to); the space occupied by added volumes; the space required for the new volume you are creating; unused available space; and the space that is overcommitted if the Pool Overcommit setting is enabled.


Volume sizes are aligned to 4.2-MB (4-MiB) boundaries. When a volume is created or expanded, the resulting size will be decreased to the nearest 4.2-MB boundary

 **NOTE:** If overcommitting the physical capacity of the system is enabled, the system will warn you via alerts (and event 462) before the pool runs out of physical storage.

 **NOTE:** You can expand the base volume of a snapshot tree, but you cannot expand any snapshots in the tree.

Rolling back virtual volumes

You can replace the data of a source volume or virtual snapshot with the data of a snapshot that was created from it using one of the methods described below.

 **CAUTION:** When you perform a rollback, the data that existed on the volume is replaced by the data on the snapshot. All data on the volume written since the snapshot was created is lost. As a precaution, create a snapshot of the volume before starting a rollback.

- Volumes table (**Provisioning > Volumes**)
- Data Protection table (**Provisioning > Volumes > Data Protection**)
- Snapshots panel (**Provisioning > Volumes > slide-over > Snapshots**)

Select the volume or snapshot, then choose **Rollback Volumes** from the drop-down list. Follow the on-screen directions to complete the action.

 **NOTE:** Volumes protected by a snapshot have a  icon under Protection Type on the Data Protection tab.

For more information see [Snapshots](#).

Creating snapshots

You can create snapshots of a volume by accessing one of the following:

- Volumes table (**Provisioning > Volumes**)—Select the volume, and choose **Add Data Protection** from the drop-down list.
- Volumes table (**Provisioning > Volumes > Data Protection > Add Data Protection**)
- Snapshots panel (**Provisioning > Volumes > slide-over > Snapshots**)

Follow the on-screen directions to complete the action. For more information see [Snapshots](#).

Resetting snapshots

You can replace the data of a standard snapshot with the current data from its parent volume by accessing one of the following:

- Volumes table (**Provisioning > Volumes**)
- Volumes table (**Provisioning > Volumes > Data Protection**)
- Snapshots panel (**Provisioning > Volumes > slide-over > Snapshots**)

Select the volume and choose **Reset Snapshot** from the drop-down list. Follow the on-screen directions to complete the action.


Copying volumes or snapshots

You can copy a volume or snapshot to a new volume by accessing one of the following:

- Volumes table (**Provisioning > Volumes**)
- Volumes table (**Provisioning > Volumes > Data Protection**)
- Overview panel (**Provisioning > Volumes > slide-over > Overview**)
- Snapshots panel (**Provisioning > Volumes > slide-over > Snapshots**)


Select the volume and choose **Copy Volume** from the drop-down list. Follow the on-screen directions to complete the action.

For more information, see [Copying volumes or snapshots](#).

 **NOTE:** You must detach the volume before copying it.

Aborting a volume copy

You can abort a volume copy operation (**Provisioning > Volumes**) by selecting the slide-over of the volume being copied.


In the Overview panel, click  next to the progress indicator. Follow the prompts to abort the operation.

Add data protection

Choosing **Add Data Protection** opens the Data Protection Configuration wizard where you are led through the process of adding data protection to a selected volume by either creating local snapshots or creating a remote replication set.

The options that display in the wizard are dependent upon how the wizard is accessed:

- **Provisioning > Volumes > menu list > Add Data Protection**
- **Provisioning > Volumes > Data Protection > Add Data Protection**
- **Provisioning > Volumes > slide-over > Snapshots**
- **Provisioning > Volumes > slide-over > Replications**—Follow the on-screen directions to complete the action.

 **NOTE:** When creating a remote replication or local snapshots, you will be prompted to set a schedule. This is the only place where you can set either a replication or snapshot schedule using the PowerVault Manager. If you do not set a schedule when prompted, you will be unable to do so using the PowerVault Manager after the replication set or snapshot is created. For information about creating a schedule using the `create schedule` CLI command, see the CLI Reference Guide.

Related information

[Replication](#)

[Snapshots](#)

[Modifying a replication set](#)

[Deleting a replication set](#)

[Initiating or scheduling a replication set](#)

[Aborting a replication set](#)

[Suspending a replication](#)

[Resuming a replication](#)

[Managing replication schedules](#)

Creating a replication set

Choosing **Add Data Protection** opens the Data Protection Configuration wizard where you are led through the process of creating a replication set for a selected volume.


The options that display in the wizard are dependent upon how the wizard is accessed:

- **Provisioning > Volumes > menu list > Add Data Protection**
- **Provisioning > Volumes > Data Protection > Add Data Protection**
- **Provisioning > Volumes > slide-over > Replications**—Follow the on-screen directions to complete the action.

NOTE: When creating a remote replication, you will be prompted to set a replication schedule. This is the only place where you can schedule a replication set using the PowerVault Manager. If you do not set a replication schedule when prompted, you will be unable to do so using the PowerVault Manager after the replication set is created. For information about creating a schedule using the `create schedule` CLI command, see the CLI Reference Guide.

Modifying a replication set

You can modify a replication set name, queue policy, snapshot history settings, and associated schedules (**Provisioning > Volumes**) by selecting the slide-over to access the Replications panel.

Click  next to the options you want to modify in the Replication Details section.

NOTE: Modifying a replication set name does not affect the schedule.

NOTE: Consider the following points before modifying a replication set:

- If you change the Snapshot Base Name while a replication is running, for the current replication it will affect the name of the snapshot on the secondary system. For that replication only, the names of the snapshots on the primary and secondary systems will differ.
- If you reduce the snapshot count setting to a value less than the current number of snapshots, the operation will fail. Thus, you must manually delete the excess snapshots before reducing the snapshot count setting. If you change this parameter while a replication is running, for the current replication it will affect only the secondary system. In this case the value can only be increased, so you might have one less expected snapshot on the primary system than on the secondary system.

Related information

[Queuing replications](#)

[Maintaining replication snapshot history](#)

Deleting a replication set

You can delete a replication set (**Provisioning > Volumes**) by selecting the slide-over to access the Replications panel.

Click **Remove Replication**. The replication set is deleted.

When you delete a replication set, all infrastructure created by the system (internal snapshots required to support replications) is also deleted. The primary and secondary volumes no longer have restrictions and function like all other base volumes and snapshots.

If you want to delete a replication set that has a replication in progress, you must first suspend, and then abort, the replication for that replication set. When a replication set is deleted, the snapshots created from the snapshot history feature will not be deleted; however, the schedule for the replication set will be deleted if one had been created. You will be able to manage those snapshots like any other snapshots.

NOTE: If the peer connection is down and there is no communication between the primary and secondary systems, use the `local-only` parameter of the `delete replication-set` CLI command on both systems to delete the replication set. For more information, see the CLI documentation.

Related information

[Aborting a replication set](#)

[Suspending a replication](#)

[Maintaining replication snapshot history](#)

Initiating or scheduling a replication

After you have created a replication set, you can copy the selected volume on the primary system to the secondary system by initiating replication (**Provisioning > Volumes**). Select the slide-over to access the Replications panel, and then click **Start Replication**.

- If a replication is not in progress, the local system begins replicating the contents of the replication set volume to the remote system and a progress bar indicates the status of the replication set.
- If a replication set is already in progress, then the outcome of this replication request depends on the Queue Policy setting specified. For more information on setting the queue policy, see [Queuing replications](#).

The first time that you initiate replication, a full copy of the allocated pages for the volume is made to the secondary system. Thereafter, the primary system only sends the contents that have changed since the last replication.

You can manually initiate replication after you create a replication set. You can create a scheduled task to automatically initiate replication from the Data Protection Configuration wizard when you create the replication set. You can initiate replications and manage replication schedules only from the primary system of a replication set.

NOTE: When creating a remote replication, you are prompted to create a replication schedule. This is the only time where you can schedule a replication set using the PowerVault Manager. If you do not set a replication schedule when prompted, you will be unable to do so in the PowerVault Manager after the replication set is created. For information about creating a schedule using the create schedule CLI command, see the CLI Reference Guide.

NOTE: If you change the time zone of the secondary system in a replication set whose primary and secondary systems are in different time zones, you must reset the secondary system to enable management interfaces to show proper time values for replication operations, such as the start, end and estimated completion replication times. To reset the time values, log into one of the controllers on the secondary system, and restart only one of Storage Controllers using either the `restart sc a` or `restart sc b` command.

If a replication operation encounters a problem, the system suspends the replication set. The replication operation will attempt to resume if it has been more than 10 minutes since the replication set was suspended. If the operation has not succeeded after six attempts using the 10-minute interval, it will switch to trying to resume if it has been over an hour and the peer connection is healthy.


NOTE: Host port evaluation is done at the start or resumption of each replication operation.

- At most, two ports will be used.
- Ports with optimized paths will be used first. Ports with unoptimized paths will be used if no optimized path exists. If only one port has an optimized path, then only that port will be used.
- Ports with both paths optimized will share the replication traffic between the two ports. The amount of replication traffic each port moves is dependent upon other operations occurring on each port.
- The replication will not use another available port until all currently used ports become unavailable.



NOTE: If a single host port loses connectivity, event 112 will be logged. Because a peer connection is likely to be associated with multiple host ports, the loss of a single host port may degrade performance but usually will not cause the peer connection to be inaccessible.

Suspending a replication

You can suspend replication operations for a specified replication set (**Provisioning > Volumes**) by selecting the slide-over to access the Replications panel.

If the replication is in progress, click  next to the progress indicator. If the replication is in an unsynchronized or ready state, click **Suspend Replication**.

You can suspend replications only from the primary system of a replication set.

When you suspend a replication set, all replications in progress are paused and no new replications are allowed to occur. The replication set will stay in a suspended state until you resume  or abort  the replication.

Resuming the replication set continues the replications that were in progress and allows new replications to occur. Aborting the replication set cancels the replication. For more information, see [Aborting a replication set](#) or [Resuming a replication](#).


If replications are attempted while the operation is suspended, such as manual or scheduled replications, the replications will fail.

Aborting a replication set

Aborting a replication set cancels the replication. You can abort running or suspended replication operations for a specified replication set (**Provisioning > Volumes**) by selecting the slide-over to access the Replications panel.


Click  next to the progress indicator.

You can abort replications only from the primary system of a replication set.

 **NOTE:** If you abort the initial replication for a replication set, the snapshot space allocated for that replication in the primary pool and the secondary pool will not be freed. To free that space, either re-run the initial replication or delete the replication set.

Resuming a replication



You can resume the replication operations of a specified suspended replication set (**Provisioning > Volumes**) by selecting the slide-over to access the Replications panel.


Click  next to the progress indicator.

When a replication set is suspended, all replications in progress are paused and no new replications are allowed to occur. When you resume replications, all paused replications are resumed and new replications are allowed to occur depending on the queue policy setting. For more information on setting the queue policy, see [Queuing replications](#).

Managing replication schedules

You can modify or delete scheduled replication tasks from the primary system (**Provisioning > Volumes**) by selecting the slide-over to access the Replications panel.

Click  in the Replication Schedules table and follow the on-screen directions to modify the schedule. Click  to delete the schedule.

 **NOTE:** This option is unavailable if a replication set schedule is not defined. See [Add data protection](#) for more information.

Working with hosts


The Hosts panel (**Provisioning > Hosts**) provides options to create hosts and host groups, display a list of existing hosts, host groups, and initiators that are a part of an existing host or host group, and display a list of all initiators.

For more information about creating hosts, see [Creating hosts](#). To perform an action on an existing host or host group, select one or more hosts, host groups, or initiators from the table and then choose an option from the drop-down list:


- [Attaching hosts to volumes](#)
- [Detaching hosts from volumes](#)
- [Removing initiators from a host](#)
- [Removing hosts from a host group](#)
- [Adding hosts to a host group](#)
- [Deleting hosts](#)
- [Deleting host groups](#)

Other actions to take on this tab include:

- [Renaming hosts](#)
- [Changing a host profile](#)
- [Renaming host groups](#)
- [Renaming initiators](#)

Click  to expand the host row to see initiator details. Select a host or initiator to perform an action from the drop-down list.

Click the host slide-over to view the Overview tab, where you can edit the name of the host and nickname of each initiator. Click the **Attached Volumes** tab to see information about attached volumes, attach a volume to the host, and to unmap volumes from the host. Follow the on-screen directions for more details.

Click the initiator slide-over  to view the Overview tab and see initiator details. Click the **Attached Volumes** tab to see information about volumes attached to the initiator.

Click the **All Initiators** tab to display a list of existing initiators on the system. To perform an action, select one or more initiators from the table and then choose an option from the drop-down list:


- [Adding initiators to a host](#)
- [Removing initiators from a host](#)


For more information about hosts and initiators, see [Initiators, hosts, and host groups](#).

Creating hosts

Click **Provisioning > Hosts > Create Hosts** to open the Create Hosts wizard to create hosts and host groups from existing initiators.

Follow the on-screen directions to create one or more new hosts and attach those hosts or host groups to initiators. Fields with a red asterisk are required. The wizard prompts you to create a new host or host group, add initiators, and create or select a volume to attach to the host or host group. All selected volumes will be attached to the newly created host.

 **NOTE:** If you have a small monitor, you may need to scroll to the bottom of the wizard to see all of the available options.

 **NOTE:** If your storage configuration has virtual pools greater than 2 PB, use host-side driver settings to increase the host I/O timeout interval (Block Device Timeout) to 80 seconds.

Related information

[Volumes and volume groups](#)

[Initiators, hosts, and host groups](#)

Attaching hosts to volumes

Attach hosts to volumes from the Hosts table (**Provisioning > Hosts**) by selecting the host and choosing **Attach Volumes** from the drop-down list or from the Attached Volumes panel (slide-over **> Attached Volumes**). Follow the on-screen directions to complete the action.

Detaching hosts from volumes

Detach hosts from volumes from the Hosts table (**Provisioning > Hosts**) by selecting the host and choosing **Detach Volumes** from the drop-down list or from the Attached Volumes panel (slide-over **> Attached Volumes**). Follow the on-screen directions to complete the action.

Removing initiators from a host

You can remove initiators from a host or host group from the Hosts table (**Provisioning > Hosts > All Initiators**) by selecting the initiator and choosing **Remove From Host** from the drop-down list. Follow the on-screen directions to complete the process.

Removing an initiator from a host will ungroup the initiator, but will not delete it or change its mapping. This action is disabled if:

- The selected initiator is the only one attached to the host. You must delete the host to free up in the initiator.
- The selected initiator is not currently attached to a host.

Removing hosts from a host group

You can remove hosts from a host group (**Provisioning > Hosts**) by selecting the host from the Hosts table and choosing **Remove From Host Group** from the drop-down list. Follow the on-screen directions to complete the process.

Removing a host from a host group will ungroup the host, but will not delete it or change its mapping. To delete a host group, see [Deleting host groups](#).

Adding hosts to a host group

You can add hosts to a new or existing host group from the Hosts table (Provisioning > Hosts) by selecting the host or host group and choosing **Add to Host Group** from the drop-down list. Follow the on-screen instructions to complete the process.

Keep the following rules in mind when adding hosts to a host group:

- The host must be attached with the same access, port, and LUN settings to the same volumes as every other host in the host group.
- A host group can contain a maximum of 256 hosts.

Deleting hosts

You can delete hosts that are not grouped (**Provisioning > Hosts**) by selecting the host from the Hosts table and choosing **Delete Host** from the drop-down list. Follow the on-screen directions to complete the process.

Deleting a host will ungroup its initiators, but they will still be visible if they are physically connected to the system. The host will detach from any attached volumes and the host device will lose access to all volume data.

Deleting host groups


You can delete host groups (**Provisioning > Hosts**) by selecting the host group from the Hosts table and choosing **Delete Host Group** from the drop-down list. Follow the on-screen directions to complete the process.

Deleting a host group will ungroup the hosts from the group but will not delete them. You will lose access to any volumes that were attached to the host group. You will retain access to any volumes that were attached to hosts in the group.

Add initiators to a host

You can add existing initiators to an existing host from the Hosts table (**Provisioning > Hosts > All Initiators**) by selecting the initiator and choosing **Add to Existing Host** from the drop-down list. Follow the on-screen directions to complete the process.


Keep the following rules in mind when adding initiators to a host:

- The initiator must be attached with the same access, port, and LUN settings to the same volumes as every other initiator in the host. An initiator must be named (nicknamed) to be added to a host; if it is not already named, a default name will be assigned. Click the initiator slide-over  to edit/add an initiator nickname.
- A host can contain a maximum of 128 initiators.

Renaming hosts

You can rename hosts from the Overview panel (**Provisioning > Hosts > Hosts and Host Groups >** slide-over). Click  next to the hostname to modify it.

Changing a host profile

You can change the profile for the initiators of hosts from the Overview panel (**Provisioning > Hosts > Hosts and Host Groups >** slide-over). Click  within the Hosts table, then choose an option from the Profile drop-down menu .

Renaming a host group

You can rename host groups from the Overview panel (**Provisioning > Hosts > Hosts and Host Groups >** slide-over). Click  next to the host group to modify it.

Renaming initiators

You can rename initiator nicknames from the Overview panel (**Provisioning > Hosts > Hosts and Host Groups >** slide-over). Click  next to the initiator name to modify it.

You can also edit an initiator nickname from (**Provisioning > All Initiators >** slide-over). Click  next to the initiator name to modify it.

Settings

Use the Settings panel to view and manage system configuration settings.

Topics:

- [Network settings](#)
- [User settings](#)
- [System settings](#)
- [Notification settings](#)
- [iSCSI host port settings](#)
- [Peer connection settings](#)

Network settings

The Network panel (**Settings > Network**) provides options for you to configure IPv4 and IPv6 network-port settings, configure a DNS server, enable or disable system management services, and view certificates.

Configuring controller network ports

The system provides concurrent support for IPv4 and IPv6 protocols. Both protocols can be set up at the same time by configuring the network parameters.

You can manually set static IP address parameters for network ports, or you can specify that IP values be set automatically, using DHCP (Dynamic Host Configuration Protocol) for IPv4 or DHCPv6 or SLAAC (Stateless address auto-configuration) for IPv6.

NOTE: SLAAC relies on Neighbor Discovery Protocol (NDP), and is the simplest way to provide an IPv6 address to a client.

If (**Settings > Network > IPv6 > (controller A|B) > Source > Auto**) is selected, the system will use an automated method-defined via the network configuration: which could be DHCPv6 or SLAAC-to auto-configure the address. The Auto setting presents a single IPv6 address. If a DHCPv6 address is available, DHCPv6 will provide the interface address; otherwise, the SLAAC address will be used.

When setting IP address values, you can choose IPv4 formatting, IPv6 formatting, or both for each controller. Additionally, you can set the addressing mode and IP address version differently for each controller and use them concurrently. For example, you could set IPv4 on controller A to Manual to enable static IP addressing, and IPv6 on controller A to Auto to enable automatic IP addressing. Given that network parameter settings are independent between the two protocols, you can set them as needed for IP addressing on controller B.

When using DHCP mode, the system obtains values for the network port IP address, subnet mask, and gateway from a DHCP server if one is available. If a DHCP server is unavailable, the system will use its default values. You must have some means of determining what addresses have been assigned, such as the list of bindings on the DHCP server. You can retrieve the DHCP assigned IP addresses either through the USB serial console login page, which lists IPv4 and IPv6; via CLI commands; or from the DHCP server list of MAC address to IP address leases. When using Auto mode, addresses are retrieved from both DHCP and SLAAC. DNS settings are also automatically retrieved from the network.

Each controller has the following factory-default IP settings:

- IP address source: Manual
- Controller A IP address: 10.0.0.2
- Controller B IP address: 10.0.0.3
- IP subnet mask: 255.255.255.0
- Gateway IP address: 10.0.0.1

NOTE: The following IP addresses are reserved for internal use by the storage system: 169.254.255.1, 169.254.255.2, 169.254.255.3, and 169.254.255.4. Because these addresses are routable, do not use them anywhere in your network.

For IPv6, when Manual mode is enabled you can enter up to four static IP addresses for each controller. When Auto is enabled, the following initial values are set and remain set until the system is able to contact a DHCPv6 and/or SLAAC server for new addresses:

- Controller A IP address: fd6e:23ce:fed3:19d1::1
- Controller B IP address: fd6e:23ce:fed3:19d1::2
- Gateway IP address: fd6e:23ce:fed3:19d1::3


 **CAUTION: Changing IP address settings can cause management hosts to lose access to the storage system after the changes are applied in the confirmation step.**

Once you set the type of controller network ports to use, you can configure domain names using the Domain Name Service (DNS). DNS accepts IPv4 and IPv6 address formats. For more information about the DNS feature, see [DNS settings](#).

Configuring DNS settings

Configure domain names using DNS (**Settings > Network > DNS**) after setting the type of controller network ports to use (IPv4 and/or IPv6).

When configuring DNS settings, note the following:

- The system supports a maximum of three DNS servers per controller.
- DNS hostnames must differ for each controller, are not case sensitive, and can have from 1 to 63 bytes. The name must start with a letter and end with a letter or digit, and can include letters, numbers, or hyphens, but no periods.
- In the DNS Servers fields for each controller, specify up to three named server addresses that are recognized within your network to be queried by the DNS resolver. The resolver will query the network in the order listed until reaching a valid destination address. Any valid setting is treated as enabling DNS resolution for the system.
- In the DNS Search Domains fields for each controller, specify up to three domain names to search when resolving hostnames that are configured in the storage system. The resolver will query the network in the order listed until finding a match.
- To reset a hostname for a controller to its default setting, click the **Reset Host Name** button for that controller.
- To clear a DNS server or search domain for a controller, click  for that setting.
- To clear all configured DNS servers and search domains for a controller, click the **Clear DNS** button for that controller. For more information about the DNS feature, see [DNS settings](#).

Enabling or disabling system-management services

You can enable or disable management interface services to limit the ways in which users and host-based management applications can access the storage system. Network management services operate outside the data path and do not affect host I/O to the system. To allow specific users to access the PowerVault Manager, CLI, or other interfaces, see User settings.

For information about setting options to control CLI behavior—to include setting its output mode—see the `set cli-parameters` command in the CLI Reference Guide.

Enable the services that you want to use to manage the storage system and disable others by choosing options on the Services panel (**Settings > Network > Services**).

- Web and API
 - **HTTPS**—Enable the use of HTTPS to provide secure access to the PowerVault Manager, the web application that is the primary interface for managing the system.
 - **HTTP**—Enable the use of HTTP to provide access to the PowerVault Manager.
- Command Line Interface
 - **SSH**—Enable the use of SSH to provide secure access to the CLI, an advanced-user interface used to manage the system and to write scripts or to run scripts. By default, SSH is enabled.
 - **Telnet**—Enable the use of Telnet to access the CLI to manage the system and to write scripts or run scripts. By default, Telnet is disabled.
 - **SSH Port**—If you enable SSH, specify the port number to use. The default is 22.
- File Transfer Protocol
 - **FTP**—A secondary interface for uploading firmware updates and downloading logs.
 - **SFTP**—A secure secondary interface for uploading firmware updates, downloading logs, and installing security certificates and keys. All data sent between the client and server will be encrypted. SFTP is enabled by default
 - **SFTP Port**—If you enable SFTP, specify the port number to use. The default is 1022.
- Other Interfaces

- **SNMP**—Enables or disables Simple Network Management Protocol (SNMP). SNMP is used for remote monitoring of the system through your network
- **SLP**—Enables or disables the Service Location Protocol (SLP) interface. SLP is a discovery protocol that enables computers and other devices to find services in a LAN without prior configuration. This system uses SLP v2.

Viewing certificate information

You can view controller certificate information in the Certificates panel (**Settings > Network > Certificates**). By default, the system generates a unique SSL certificate for each controller. For the strongest security, replace the default system-generated certificate with a certificate issued from a trusted certificate authority.

The Certificates panel shows information for the active SSL certificates that are stored on the system for each controller. Controller A and Controller B tabs contain unformatted certificate text for each of the corresponding controllers. The panel also shows one of the following status values as well as the creation date for each certificate:

- **Customer-supplied**—Indicates that the controller is using a certificate that you have uploaded.
- **System-generated**—Indicates that the controller is using an active certificate and key that were created by the controller.
- **Unknown status**—Indicates that the controller module certificate cannot be read. This situation occurs most often when a controller is restarting or certificate replacement is still in progress, or you have selected the tab for a partner controller in a single-controller system

You can use your own certificates by uploading them through SFTP or by using the contents parameter of the `create certificate` CLI command to create certificates with your own unique certificate content. For a new certificate to take effect, you must restart the affected Management Controller. To restart a controller, select **Maintenance > Hardware > Rear View > Enclosure Actions > Restart/Shutdown System > Restart MC** and follow the on-screen instructions.

To verify that the certificate replacement was successful and the controller is using the certificate that you have supplied, make sure the certificate status is customer-supplied, the creation date is correct, and the certificate content is the expected text.

For details pertaining to supported certificate criteria, see [Installing a security certificate](#).

User settings


The Users panel (**Settings > Users**) provides options for you to manage local users and SNMPv3 users. Options on this panel let you add, modify, and delete users; set user permissions; and set system preferences based on individual user profiles.

Managing local users


The Local Users panel (**Settings > Users > Local**) provides options to add new users and modify system permissions for existing users. The first user that completes the onboarding process during system setup will have the manage role. A user with the manage role can add up to nine additional users (SNMPv3 users count towards this limit), modify any user, and delete any user other than the current user.

Users assigned a standard or monitor role can change their own username, password, language and temperature preference, and timeout setting. Standard and monitor users cannot change their access to user interfaces or roles, and they cannot change the settings of other users.

Manage and standard users can both access one or more of the following management interfaces: the PowerVault Manager, CLI or FTP and SFTP. Monitor users can only access the PowerVault Manager and the CLI management interfaces.

 **NOTE:** To secure the system, each user should have a unique username and password.

Local user options

The following options are available to users with a manage or standard role when adding or modifying users. To add new users, click **Add New Users**, and to modify users click .

- **Username**—A username is case sensitive and can have a maximum of 29 bytes. The name cannot already exist in the system, include spaces, or include any of the following: " , < \
- **Password**—A password is case sensitive and can have from 8 to 32 printable characters. If the password contains only ASCII characters, it must contain at least one uppercase character, one lowercase character, one numeric character, and one non-alphanumeric character. A password can include UTF-8 characters except for the following: a space or " ' , < > \

- **Language**—Select a display language for the user. The default is English. Installed language sets include Chinese-Simplified, English, French, German, Japanese, Korean, and Spanish. The locale determines the character used for the decimal (radix) point. The locale setting is determined by the Language setting, which can be accessed by selecting the pencil icon for any user in the table.
- **Temperature Preference**—Select whether to use the Celsius or Fahrenheit scale for display of temperatures. The default is Celsius.
- **Timeout**—Select the amount of time that the user session can be idle before the user is automatically signed out (from 2 to 720 minutes). The default is 30 minutes.


The following options are available to users with a manage role when adding or modifying users:

- **Interfaces**—Select one or more of the following interfaces:
 - **WBI**. Enables access to the PowerVault Manager.
 - **CLI**. Enables access to the command-line interface.
 - **FTP**. Enables access to the FTP interface or the SFTP interface, which can be used instead of the PowerVault Manager to install firmware updates and to download logs.
- **Roles**—Select one or more of the following roles:
 - **Manage**. Enables the user to change system settings.
 - **Monitor**. Enables the user to view but not change system status and settings.

Managing SNMPv3 users

The SNMPv3 Users panel (**Settings > Users > SNMPv3**) provides options to create SNMPv3 users who can either access the Management Information Base (MIB) or receive trap notifications. SNMPv3 users manage SNMPv3 security features, such as authentication and encryption.

For information about the MIB, see [SNMP reference](#). When creating an SNMPv3 user, the system verifies whether the SNMP setting is enabled (**Settings > Network > Services**). If it is not enabled, a warning informs that the SNMP setting will be auto-enabled as the SNMPv3 user is created on the storage system.

 **NOTE:** The Engine ID is used to uniquely identify SNMPv3 entities. The Engine ID will be generated from the controller MAC address.

SNMPv3 user options

The following options are available to SNMPv3:

- **Username**—A username is case sensitive and can have a maximum of 29 bytes. The name cannot already exist in the system, include spaces, or include any of the following: " , < \
- **Password**—A password is case sensitive and can have from 8 to 32 printable characters. If the password contains only ASCII characters, it must contain at least one uppercase character, one lowercase character, one numeric character, and one non-alphanumeric character. A password can include UTF-8 characters except for the following: a space or " ' , < > \
- **Authentication Type**—Select whether to use MD5 or SHA (SHA-1) authentication, or no authentication. If authentication is enabled, the password set in the Password and Confirm Password fields must include a minimum of 8 characters and follow the other SNMPv3 privacy password rules.
- **Privacy Type**—Select whether to use DES or AES encryption, or no encryption. To use encryption you must also set a privacy password and enable authentication.
- **Privacy Password**—If the privacy type is set to use encryption, specify an encryption password. This password is case sensitive and can have from 8 to 32 printable characters. If the password contains only ASCII characters, then it must contain at least one uppercase character, one lowercase character, one numeric character, and one non-alphabetic character. A password can include UTF-8 characters except for the following: a space or " ' , < > \
- **Trap Host Address**—Specify the network address of the host system that will receive SNMP traps. The value can be an IPv4 address, IPv6 address, or FQDN.

System settings

The System panel (**Settings > System**) provides options for you to set system identification information, set the system date and time, secure the system using FDE, and set system properties.

Setting system identification information

The Identification panel (**Settings > System > Identification**) provides options for you to specify information to identify the system.

Enter the name of the system, the name of the person or team that manages the system, the location of the system, and any additional information about the system's use or configuration. The system name is shown in the PowerVault Manager browser title bar or tab and is included in notification emails. All of the information is included in system debug logs for reference by service personnel.

Setting the date and time


Set the date and time (**Settings > System > Date and Time**) so that entries in system logs and notifications have correct time stamps. The banner displays the system date and time in the format `<year>-<month>-<day> <hour>:<minutes>:<seconds>`.


Access the Date and Time panel by clicking on the date and time displayed in the banner or by clicking **Settings > System > Date and Time**.

You can set the date and time manually or configure the system to use Network Time Protocol (NTP) to obtain date and time from an available network-attached server. Using NTP allows multiple storage devices, hosts, log files, and so on to be synchronized. The NTP server address value can be an IPv4 address, IPv6 address, or FQDN. If NTP is enabled but no NTP server is present, the date and time are maintained as if NTP was not enabled.

NTP server time is provided in the UTC time scale, which provides several benefits:

- To synchronize the times and logs between storage devices installed in multiple time zones, set all the storage devices to use UTC.
- To use the local time for a storage device, set its time zone offset.
- If a time server can provide local time rather than UTC, configure the storage devices to use that time server, with no further time adjustment.

 **NOTE:** Whether NTP is enabled or disabled, the storage system does not automatically make time adjustments for Daylight Saving Time. You must make such adjustments manually by changing the time zone offset.

 **NOTE:** If you change the time zone of the secondary system in a replication set whose primary and secondary systems are in different time zones, you must restart the secondary system to enable management interfaces to show proper time values for replication operations, such as the start, end and estimated completion replication times.

Securing the system with FDE

The Full Disk Encryption panel (**Settings > System > Security**) provides options for you to enable FDE protection to secure all of the user data on a storage system. To secure the system, all disks must be FDE-capable.

The Security panel indicates whether or not the system is secured. To secure the system, click **Secure System** to enter a passphrase that will enable the system to write to and read from FDE-capable disks. The system uses the passphrase to generate a lock key ID that is used to secure the FDE-capable disks. If the system is unable to interpret the lock key on the FDE-capable disk, the encrypted data on the disk is inaccessible.


 **CAUTION:** Be sure to record the passphrase as it cannot be recovered if lost.

Once the passphrase is entered and confirmed, the System Status will indicate that the system is secured and the Security panel will provide options to:

- Change the passphrase. Enter a new passphrase.
- Lock the system for transport. Lock down the disks in preparation for transport. Use this option when the system will not be under your physical control. After the system has been transported and powered up, the system and disks will enter the secured, locked state; disks will be in the UNUSABLE state:

- For virtual configurations, pools and disk groups will be unavailable. To restore access to encrypted data, enter the passphrase for the system's lock key ID. Disk groups will be dequarantined, pool health will be restored, and volumes will become accessible
- For linear configurations, disk-groups will be unavailable. To restore access to encrypted data, enter the passphrase for the system's lock key ID. Disk groups will be dequarantined, health will be restored, and volumes will become accessible

Lock keys are generated from the passphrase and manage locking and unlocking the FDE-capable disks in the system. Clearing the lock keys and power cycling the system denies access to encrypted data on the disks. If the lock keys are cleared while the system is secured, the system will enter the FDE lock-ready state, in preparation for the system being powered down and transported.

 **CAUTION: Do not change FDE configuration settings while running I/O. Temporary data unavailability may result, and the proper setting of lock keys from the passphrase could potentially be impacted.**


Repurposing secured disks

Click **Repurpose Secured Disks** to repurpose a disk that is no longer part of a disk group.

Repurposing a disk resets the encryption key on the disk, effectively deleting all data on the disk. After a disk is repurposed in a secured system, the disk is secured using the system lock key ID and the new encryption key on the disk, making the disk usable to the system.

Repurposing a disk in an unsecured system removes all associated lock keys and makes that disk available to any system.

This action is also available via **Maintenance > Hardware**. Once you have selected the secured disks from the graphic view, select Repurpose Secured Disks from the Related Health Actions dropdown and follow on-screen instructions.

 **CAUTION: Repurposing a disk changes the encryption key on the disk and effectively deletes all data on the disk. Repurpose a disk only if you no longer need the data on the disk.**

The Repurpose Secured Disks action is not permitted when the system is in the secured, locked state

For more information about using FDE, see [Full disk encryption](#).

Setting import lock key IDs

You can set the passphrase associated with an import lock key to unlock FDE-secured disks that are inserted into the storage system from a different secure system using either of two methods.

- **Settings > System > Security > Full Disk Encryption > Import Secured Disks**. Use tool tips and follow on-screen instructions to enter the passphrase and integrate locked disks into the system.
- **Maintenance > Hardware**. Once you have selected the secured disks from the graphic view, select **Import Secured Disks** from the **Related Health Actions** dropdown, and follow on-screen instructions to enter the passphrase and integrate locked disks into the system.

If the correct passphrase is not entered, the storage system cannot access data on the disk. After importing disks into the system, they will be associated with the system lock key ID, and data will no longer be accessible using the import lock key. This effectively transfers security to the local system passphrase.

Setting system properties

Use the System Properties panel to change system cache properties, disk properties, scrub properties, managed logs properties, and firmware properties.

For more information about setting advanced system configuration parameters, see the `set advanced-settings` CLI command within the CLI Reference Guide.

Setting system cache properties

The Cache Properties panel (**Settings > System > Properties > Cache Properties**) lets you set the synchronize-cache mode, missing LUN response, host control of the system's write-back cache setting, and auto-write-through cache triggers.

If you are experiencing performance issues, verify that the cache properties are set to optimize system performance. See the help tips in the panel for specific information regarding each option.

NOTE: You can change the cache parameters for a specific volume from the Overview tab of the Volumes table (**Provisioning > Volumes >** slide-over). For more information on performance cache options, see the `set volume-cache-parameters` CLI command in the CLI Reference Guide.

Setting system disk properties

The Disk Properties panel (**Settings > System > Properties > Disk Properties**) provides options to do the following:

- Enable disk monitoring and failure analysis (SMART)
- Change polling frequency to alert you to temperature changes, power supply and fan status, and the presence or absence of disks
- Enable dynamic spare capability
- Enable disk spin down (DSD)

See the help tips in the panel for specific information regarding each option. For more information about dynamic spares, see [Spares](#).

Setting system scrub properties

The Scrub Properties panel (**Settings > System > Properties > Scrub Properties**) provides options to enable scrub operations to inspect and fix errors found on disks and in disk groups. Dell recommends that both of these options be enabled. See the help tips in the panel for specific information regarding each option.

NOTE: If scrub options are disabled, you can still scrub a selected disk group. See [Disk group scrub](#).

Setting system managed logs properties

Enabling the Managed Logs feature (**Settings > System > Properties > Managed Logs Properties**) transfers system log files containing diagnostic data to an external log collection system for retention.

For more information see [Managed logs](#).

Entering an email address in the Log Destination Email text box will enable the system to attach log files to managed-log email notifications sent by the log collection system. See the help tips in the panel for specific information regarding each option.

Setting partner firmware update

When Partner Firmware Update is enabled (**Settings > System > Properties > Firmware Properties**), firmware on the partner controller is automatically updated when firmware on the primary controller is updated.

NOTE: PFU should be enabled. Disable this option only if told to do so by a service technician.

Notification settings

The Notifications panel (**Settings > Notifications**) provides options to send system alert notifications to users through email, SNMP trap hosts, or a remote syslog server.

For more information about alerts, see [Alerts panel](#).


NOTE: You should enable at least one notification service to monitor the system.

Email notifications

You can choose to be notified by email when system alerts occur.

Alert notifications can be sent to a maximum of three email addresses. Weekly alerts concerning system health issues will also be sent until corrective action has been taken and the system health value has returned to OK.

Enter information in the text boxes to receive alert notifications. For details about panel options, see the on-screen tool tips. For information about SMTP notification parameters for events and managed logs, see the `set email-parameters` command in the CLI Reference Guide.

 **NOTE:** If the mail server is not on the local network, make sure that the gateway IP address was set in [Configuring controller network ports](#).

SNMP notifications

The SNMP panel provides options to send alert notifications to SNMP trap hosts.

You must enable SNMP for the system to send alert notifications to SNMP users. Enter information in the text boxes to receive alert notifications. For details about panel options, see the on-screen tool tips. See [Enabling or disabling system-management services](#) for more information.

Syslog notifications

The Syslog panel lets you set remote syslog notifications to allow alerts to be logged by the syslog of a specified host computer.

Syslog is a protocol for sending alert messages across an IP network to a logging server. This feature supports User Datagram Protocol (UDP), but not Transmission Control Protocol (TCP). For details about panel options, see the on-screen tool tips.

iSCSI host port settings


Use the options in the iSCSI Host Ports panel (**Settings > iSCSI > Host Ports**) to reset host links and to change the iSCSI IP address, netmask, and gateway for each port on each controller.

The panel includes the following options:

- Reset Host Links. Making a configuration or cabling change on a host might cause the storage system to stop accepting I/O requests from that host. For example, this problem can occur after moving host cables from one HBA to another on the host. To remedy this, you might need to reset controller host links (channels).
- IP Address. IPv4 or IPv6, the port IP address. For corresponding ports in each controller, assign one port to one subnet and the other port to a second subnet. Ensure that each iSCSI host port in the storage system is assigned a different IP address. For example, in a system using IPv4:
 - Controller A port 1: 10.10.10.100
 - Controller A port 2: 10.11.10.120
 - Controller B port 1: 10.10.10.110
 - Controller B port 2: 10.11.10.130

 **CAUTION:** Changing IP address settings can cause data hosts to lose access to the storage system.

- Netmask. For IPv4, enter the subnet mask for the assigned port IP address. The default is 255.255.255.0.
- Gateway. For IPv4, enter the gateway IP address for the assigned port IP address. The default is 0.0.0.0.
- Default Router. For IPv6, enter the default router for the assigned port IP address. If the gateway was set for IPv4 and then ports were switched to IPv6, the default is :: IPv4-address. Otherwise, the default is :: (the short form of all zeroes).

 **NOTE:** For information about setting host parameters, see the CLI Reference Guide.

Configuring iSCSI CHAP settings

For iSCSI, you can use Challenge Handshake Authentication Protocol (CHAP) to perform authentication between the initiator and target of a login request. To perform this authentication, a database of CHAP records must exist on the initiator and target. Each CHAP record can specify one name-secret pair to authenticate the initiator only (one-way CHAP) or two pairs to authenticate both the initiator and the target (mutual CHAP). For a login request from an iSCSI host to a controller iSCSI port, the host is the initiator and the controller port is the target.

During onboarding, you are prompted to enable CHAP, add new records, and edit and delete previously defined records. When CHAP is enabled and the storage system is the recipient of a login request from a known originator (initiator), the system will request a known secret. If the originator supplies the secret, the connection will be allowed.

To enable or disable CHAP configuration settings after onboarding is complete, check or uncheck the CHAP Authentication box (**Settings > iSCSI > Configuration**).

Regardless of whether CHAP is enabled, you can add, delete, or edit CHAP records (**Settings > iSCSI > CHAP**).

 **CAUTION: Editing or deleting CHAP records may disrupt connectivity to the host using that record.**

Special considerations apply when CHAP is used in a system with a peer connection, which is used in replication. In a peer connection, a storage system can act as the originator or recipient of a login request. If the originator has CHAP enabled-but the recipient does not-the originator is able to modify the peer connection to enable CHAP on the recipient. Provided the two systems have CHAP records for one another-and share the same secret-the recipient is able to authenticate the peer connection.

Related information


[Peer connections](#)

[CHAP and replication](#)

[Initiators, hosts, and host groups](#)

Changing iSCSI configuration settings

The iSCSI configuration panel (**Settings > iSCSI > Configuration**) provides options for you to view the system's configuration, or to alter it if the network configuration was modified. Panel options include the ability to change the IP version and to enable/disable jumbo frames to allow for larger data transfers, CHAP, and iSNS. IPv4 uses 32-bit addresses. IPv6 uses 128-bit addresses

 **CAUTION: Use extreme caution if making iSCSI configuration changes after onboarding. Modifications will disrupt connectivity to the host and disconnect the system from the network.**


Peer connection settings

The Peer Connections panel (**Settings > Peer Connections**) provides options to query a peer connection and to modify and delete peer connections.


Querying peer connections

You can query a peer connection to view information about systems you might use in a peer connection before creating the peer connection, or to view information about systems currently in a peer connection before modifying or deleting the peer connection. To query a system, specify a peer system IP address, then choose **Query Peer Connection**.

Modifying peer connection settings


Click  in the Current Peer Connections section of the panel to change the name of a current peer connection or the port address of the remote system from either the local system or the remote system without changing the peer connection type or local port settings. For example, you could configure a peer connection and then move one of the peers to a different network.

Changing the peer connection name will not affect the network connection so any running replications will not be interrupted.

 **NOTE:** Changing the remote port address will modify the network connection, which is permitted only if no replications are running and new replications are prevented from running. For the peer connection, abort any running replications and either suspend its replication sets or make sure its network connection is offline. After you have modified the peer connection, you can resume replication sets.

If CHAP is enabled on one system within a peer connection, be sure that CHAP is configured properly on the corresponding peer system before initiating this operation. For more information about configuring CHAP, see [CHAP and replication](#).

Deleting a peer connection

Click  in the Current Peer Connections section of the panel to delete a peer connection if there are no replication sets that belong to the peer connection. If there are replication sets that belong to the peer connection, you must delete them before you can delete the peer connection.

For more information see [Deleting a replication set](#).

- NOTE:** If the peer connection is down and there is no communication between the primary and secondary systems, use the `local-only` parameter of the `delete replication-set` CLI command to delete the replication set.
- NOTE:** If CHAP is enabled on one system within a peer connection, be sure that CHAP is configured properly on the corresponding peer system before initiating this operation. For more information about configuring CHAP, see [CHAP and replication](#).

Maintenance

Use the Maintenance panel to manage the system's storage configuration, hardware, and firmware. You can also view information about the storage system and perform support-related actions.

Topics:

- [Storage panel](#)
- [Hardware panel](#)
- [Firmware panel](#)
- [About panel](#)
- [Support panel](#)

Storage panel

If storage is not yet configured and provisioned for the system, this panel provides configuration options to configure your storage. Configuration options include the storage type, setting up SupportAssist, and configuring storage.

To initially configure storage, select the storage type using either the **Virtual** or **Linear** radio button. A tooltip describing these two settings is provided on the Storage panel.

If you selected Virtual as the storage type, you can select **Auto Storage Setup** to easily configure storage using the wizard. If you selected Linear as the storage type, the **Auto Storage Setup** option is not available.

The **Auto Storage Setup** button (**Maintenance > Storage**) opens the Auto Storage Setup dialog box. This dialog box provides information for the selected storage type, including the configuration suggestions.

For virtual pools, the dialog is divided into discrete sections:

- Disk health with Rescan link
- Disk groups table (with applicable configuration settings)
- Spares table
- Unused disks (if any)

The **Apply Configuration** and **Cancel** actions provide control over the automated configuration process.

For linear storage, a single storage pool contains all disk groups. The storage system provides a Pools container that includes disk groups and a Spares container.

The Storage panel (**Maintenance > Storage**) shows the system storage configuration, including pools, disk groups, and spares (non-ADAPT disk groups only) and enables you to change the configuration. See also [Pools](#) and [Spares](#).

A storage system can have one virtual pool per controller module. This panel shows a Pool table for each pool, and a Spares table. In this panel you can:

- View information about a pool.
- Change pool settings.
- View information about disk groups in a pool.
- Add a disk group to a pool.
- Rename a disk group.
- Delete a disk group from a pool.
- Expand an ADAPT disk group.
- Verify a disk group.
- Scrub a disk group.
- View information about disks in a group.
- View information about volumes that are attached to hosts.

Viewing information about a pool

If a pool contains at least one disk group, the Pool table shows this basic information:

- Size (total capacity)
- Health
- Available (unused capacity)
- Overcommit size

Expand the pool row to see this additional information:

- Whether overcommit is enabled
- Whether the pool is overcommitted
- Low threshold, medium threshold, and high threshold values
- Pool serial number

Changing pool settings

To change pool settings, in the pool row click . For more information about each setting, click .

Viewing information about volumes for each disk group in a pool

To see information about disk groups in a pool, in a Pool table expand the Disk Groups row.

For each disk group, the Disk Groups table shows this basic information:

- Name
- Controller (owning controller)
- Level (disk-protection level)
- Health
- Number of disks
- Size
- Job (type and percent of progress)

In the disk group slide-over panel, the Overview tab shows this information:

- The progress of any current job on the disk group
- Disk group name
- Serial number
- Chunk size
- Owner (preferred and current)
- Sector format
- Creation date
- Minimum disk size
- Active disk spin down:
 - Size
 - Free
- Protection level
- Number of disks
- Target spare capacity (ADAPT)
- Actual spare capacity (ADAPT)
- Status
- Stripe width ADAPT (16+2 or 8+2 options)


In the disk group slide-over panel, the Disks tab shows information about each disk. Disk location is shown in the format `<enclosure-number>.<disk-slot-number>`.

Adding a disk group to a pool

In the **Maintenance > Storage** panel, in the pool where you want to add the disk group, click **Add Disk Group** and follow the on-screen instructions.

For more information about available protection levels, see [RAID levels](#).

Renaming a disk group

In the **Maintenance > Storage** panel, locate the disk group to rename, display its slide-over panel, click  and follow the on-screen directions.

Deleting a disk group from a pool


In the **Maintenance > Storage** panel, locate the disk group to delete, click , and follow the on-screen directions

Expanding an ADAPT disk group

In the **Maintenance > Storage** panel, locate the disk group to verify, display its slide-over panel, click **Expand Disk Group**, and follow the on-screen directions.

Scrubbing a disk group

In the **Maintenance > Storage** panel, locate the disk group to scrub, display its slide-over panel, click **Scrub Disk Group**, and follow the on-screen directions.

To cancel a scrub click .

Managing spares

The Spares table (**Maintenance > Storage**) displays a list of current spares, and lets you add and remove global spares for virtual and linear disk groups, and dedicated spares for linear disk groups. The options that display in the panel are dependent upon the type of disk group selected.

Hardware panel

The Hardware panel (**Maintenance > Hardware**) shows the system hardware configuration.

The panel has three sections:

- The top section shows basic information about each enclosure: ID, Rack number, Rack position, Disk slots (used and total).
- For the selected enclosure, the middle section shows a front view or rear view of the position of components in the enclosure:
 - The front view of the 2U12 or the 2U24 shows disks with slot numbering that are accessible from the enclosure front panel.
 - The front view of the 5U84 shows disks with slot numbering for the two drawers that are accessible from the enclosure front panel.
- The dropdown beneath any of these enclosure views is context-sensitive. Its text label shows the enclosure by default, and the dropdown provides actions pertaining to the enclosure. If you select a component within the view, the dropdown text label updates to show that component, and provides component-specific actions (if available).
- For the selected enclosure or component, the bottom section shows additional information. The bottom left pane provides device details. The bottom right pane shows the health and available actions pertaining to either the enclosure or the selected component.

This table lists available actions for a given device.

Table 11. Actions available for a given device

View	Device	Available actions
Front or Rear	Enclosure	<ul style="list-style-type: none"> Restart/shutdown system Rescan all disks Turn on locator LED
Front or Rear	Disk (healthy)	Turn on locator LED
	Disk (leftover)	<ul style="list-style-type: none"> Turn on locator LED Clear disk metadata
Rear	Power supply	None
	Controller module (CM)	Turn on locator LED
	Host port	Rest host port
	Network port	None

This table lists the information shown for a given device.

Table 12. Device information displayed

Device	Information shown
Enclosures	Enclosure ID, Locator LED On/Off button, Status, Vendor, Model, Disk count, WWN, Midplane serial number, Revision, Part number, Manufacturing date, Manufacturing location, Midplane type, Enclosure power (watts), PCIE 2-capable, EMP A Revision, EMP B Revision, EMP A Bus ID, EMP B Bus ID, EMP A Target ID, EMP B Target ID
Disk module	Location, Locator LED On/Off button, LED status, Serial number, Vendor, Model, Revision, Description, Usage, Current job, Supports unmap, SMART, R/min (RPM), Size, Sector format, Transfer rate, Single pathed, Recon state, Copyback state, Disk spin count down, Temperature, Status, Power on hours, FDE state, FDE lock key
Power supply	Status, Vendor, Model, Serial number, Revision, Location, Part number, Manufacturing date, Manufacturing location. The two power supply units (PSU), numbered 0-1, reside in the PSU slots accessed from the rear of the enclosure.
Power cooling module	Status, Vendor, Model, Serial number, Revision, Location, Part number, Manufacturing date, Manufacturing location. The two power cooling modules (PCM), numbered 0-1, reside in the PCM slots accessed from the rear of the enclosure.
Controller module (CM)	Controller ID, IP Address, Description, Status, Model, Serial number, System cache memory, Revision, CPLD version, Storage Controller version, Storage Controller CPU type, Part number, Position, Hardware version, Manufacturing date, Manufacturing location. The two CMs, labeled as Controller A B, reside in the CM slots, and are accessed from the rear of the enclosure.
FC host port	Name, Port type, Status, Topology, Configured speed, Actual speed, Primary loop ID, Target ID, SFP status, Part number, Supported speeds (8G, 16G, 32G)
iSCSI host port	Name, Port type, Status, Gateway, Netmask, MAC address, IP address, IP version, ID, SFP status, Part number, Configured speed, Actual speed, 10G Compliance, Cable length, Cable technology, Ethernet compliance
SAS host port	Name, Port type, Status, Actual speed, Topology, Expected lanes, Active lanes, Disabled lanes, ID
Network port	ID, IPv4 address mode, IPv4 address, IPv4 network mask, IPv4 gateway, MAC address, IPv6 auto config, IPv6 gateway, IPv6 auto address, IPv6 manual address (1 to 4)
Expansion port	Enclosure ID, Controller ID, Name, Status
Fan module	Module name, Location, Module status (OK/not OK), Fan name (2 fans), Fan status (Up/Down), Fan speed. The five fan modules, numbered 0-4, reside in the fans slots accessed from the rear of the enclosure.

Firmware panel


The Firmware panel (**Maintenance > Firmware**) shows information about system and disk firmware versions, and enables you to perform firmware updates.

The system can store multiple system firmware bundles, including:

- **Factory firmware**—The original firmware bundle for recovery purposes or a copy of later firmware if downgrades are prohibited to the original firmware.
- **Active firmware**—The firmware bundle that is activated and in use.
- **Installed / Not Active firmware**—Another firmware bundle that is installed and available to be activated. This bundle may be a newer bundle or an older, previously active bundle.

In this panel you can:

- View information about the current firmware bundle.
- View whether the Partner Firmware Update option is enabled.
- View information about installed and active system firmware bundles.
- Install a new firmware bundle.
- Activate an installed firmware bundle.
- View information about current disk firmware and available updates.

 **NOTE:** To aid successful installation and activation of system firmware be sure to read the on-screen directions.

See also [Best practices for updating firmware](#).

Viewing information about installed and active system firmware bundles

The System tab shows this basic information about each installed bundle version:

- Bundle version
- Build date
- Status

The expanded view shows this bundle-component version information:

- GEM version (GEM package version)
- MC firmware (Management Controller)
- MC loader
- MC OS version
- CPLD revision (Complex Programmable Logic Device)
- ASIC controller version
- SC firmware (Storage Controller)


Updating system firmware

Before performing a firmware update, see [Best practices for updating firmware](#). Both controllers must run the same firmware version. Storage systems with peer connections should run the same or compatible firmware versions.

Partner firmware update (PFU) is enabled on the system by default. With PFU enabled, when you update firmware on one controller or replace a controller, the system automatically updates the partner controller. Disable PFU only if requested by a service technician. Use the PowerVault Manager or CLI to change the PFU setting.

For a dual-controller system, the partner firmware update (PFU) setting (**Settings > System > Properties > Firmware Properties**) controls how updates impact the partner controller:

- **Automatic**—PFU is enabled (the default). When you update firmware, the firmware is automatically copied to and activated on the partner controller first, and then activated on the current controller. PFU provides for updating expansion module firmware in similar fashion.

 **NOTE:** Dell recommends enabling the PFU option for controller firmware updates. PFU is enabled by default and should remain enabled. Disable this option only if instructed to do so by a qualified service technician.

- **Manual**—PFU is disabled. When you update controller module or enclosure IOM firmware on one controller, you must log in to the partner controller and manually perform the same updates.

When a controller module is installed into an enclosure at the factory, the enclosure midplane serial number and firmware-update timestamp are recorded for each firmware component in controller flash memory. This information is not erased when the configuration is changed or is reset to defaults. These two pieces of data are not present in controller modules that are not factory-installed and are used as replacements.

Updating controller firmware with the PFU option enabled ensures that the same firmware version is installed in both controller modules. PFU uses the following algorithm to determine which controller module updates its partner:

- If both controllers are running the same firmware version, no change is made.
- The controller that is installed first sends its configuration and settings to the partner controller. Similarly, if a controller is replaced, it receives configuration information from the partner controller. In both cases, the PFU setting determines subsequent firmware update behavior for both controllers.
- If both controllers were already installed in the system, the controller with firmware installed first sends its configuration and settings to the partner controller.
- If both controllers are newly installed, controller A is transferred to controller B.

To install a firmware bundle:

1. Download the firmware from Dell.com.
2. Go to **Maintenance > Firmware > System**.
3. Click, **Browse For File** and browse to the downloaded firmware bundle.
4. Follow the on-screen directions to install the firmware.

To activate a firmware bundle:

1. Go to **Maintenance > Firmware > System** and click its **Activate this Version** link to display the Activate Firmware dialog.
2. Follow the on-screen directions to start the activation.

As part of the activation process the system performs these steps: check bundle integrity, check system health, update firmware on the partner controller module, restart the partner controller module, update firmware on the local controller module, and restart the local controller module.

After the local controller module has restarted, the PowerVault Manager login screen is displayed. Log back in and go to the **Maintenance > Firmware** panel and verify that the new firmware is active on the system. An alert is also generated to inform you that the firmware has been upgraded.

If firmware activation fails, go to **Maintenance > Support > Collect Logs** and enter the necessary fields and collect the logs. The logs are needed for any support request that is generated by this failure.

 **NOTE:** Consider the following points before updating system firmware:

- A firmware update typically takes 5 minutes for a controller with current CPLD firmware, or up to 20 minutes for a controller with downlevel CPLD firmware. Expand the firmware row to view the CPLD version (**Maintenance > Firmware**).
- If the Storage Controller cannot be updated, the update operation is canceled. Verify that you specified the correct firmware file and repeat the update. Run the `check firmware-upgrade-health` CLI command to determine if any problems must be resolved before attempting to update the firmware. If this problem persists, contact technical support.
- When the firmware update on the local controller is complete, the Management Controller restarts. Until the restart is complete, sign-in pages say that the system is unavailable. When this message is cleared, you may sign in again.
- If PFU is enabled, the amount of time that is required for updating both controllers is less than 10 minutes.
- If PFU is enabled, after the firmware update has completed on both controllers, check the system health. If the system health is Degraded check the health reason. If the health indicates that the firmware version is incorrect, verify that you specified the correct firmware file and repeat the update. If this problem persists, contact technical support.
- If PFU is disabled, you must locate, install, and activate compatible firmware on the replacement expansion module.

Updating disk firmware

The Disks tab shows information about each disk in the system and its firmware, and whether firmware updates are available.

Disk location is shown in the format `<enclosure-number>.<disk-slot-number>`.

The process to update disk firmware is to upload a firmware file from the manufacturer, select the disks to update, and then apply the update to those disks. A dual-ported disk drive can be updated from either controller. For details, follow the onscreen directions and [Best practices for updating firmware](#).

CAUTION: Before updating disk firmware, stop I/O to the storage system. During the update all volumes will be temporarily inaccessible to hosts. If I/O is not stopped, mapped hosts will report I/O errors. Volume access is restored after the update completes.

Best practices for updating firmware

- In the Alerts panel on the dashboard, verify that the system health is OK. If the system health is not OK, expand the view to see the active health alerts and resolve all problems before you update firmware. For information about Active Alerts, see [Alerts panel](#)
- Run the `check firmware-upgrade-health` CLI command before upgrading firmware. This command performs a series of health checks to determine whether any conditions exist that must be resolved before upgrading firmware. Any conditions that are detected are listed with their potential risks. For information about this command, see the CLI Reference Guide
- If any unwritten cache data is present, firmware update will not proceed. Before you can update firmware, unwritten data must be removed from cache. See information about the `clear cache` command in the CLI Reference Guide.

CAUTION: Removing unwritten data may result in data loss. Contact technical support for assistance.

- If a disk group is quarantined, resolve the problem that is causing it to be quarantined before updating firmware.
- CPLD revision (Complex Programmable Logic Device)
- To ensure success of an online update, select a period of low I/O activity. This helps the update to complete as quickly as possible and avoids disruption to host and applications due to timeouts. Attempting to update a storage system that is processing a large, I/O-intensive batch job may cause hosts to lose connectivity with the storage system.
- Confirm PFU is enabled by clicking **Settings > System > Properties > Firmware Properties**.
- Do not perform a power cycle or controller restart during a firmware update. If the update is interrupted or there is a power failure, the module might become inoperative. If this occurs, contact technical support. The module might need to be returned to the factory for reprogramming.

About panel

The About panel (**Maintenance > About**) shows information about the system, its hardware, and its storage configuration.

General system information includes:

- System name
- System contact
- System location
- System information
- Product brand
- Product ID
- Platform Type
- Controller A firmware
- Controller B firmware
- Enclosure <ID>: Displays the service tag identifier for the indicated enclosure, if set

Hardware information displays the Enclosure number, SKU Part Number, and SKU Serial Number for each enclosure (disks excluded). Expand the table to see the following for each customer FRU in each enclosure:

- FRU name
- Description
- Part number
- Serial number
- Configuration serial number
- Location

Storage information includes the following for each pooldisk group:

- Disk group name

- Pool ID
- Total size
- Available size
- Snap size
- Overcommit state
- Disk groups
- Volumes
- Sector format
- Health

i **NOTE:** The actual storage information settings displayed are dependent upon the storage type selected for the system's configuration.

Support panel

Use the Support panel (**Maintenance > Support**) to perform these support-related actions:

- Enable SupportAssist and customize the information you send back:
 - Choose or change your connection type
 - Enter or update contact information for your Primary and Secondary Contact
 - Access additional advanced SupportAssist options
- Collect logs
- View the system event history
- View controller module audit logs For details, see the on-screen directions.

Configuring SupportAssist

SupportAssist must be enabled to send configuration and diagnostic information from an ME5 Series storage system to technical support. Depending on your service contract, SupportAssist also automates support request creation for issues that are detected on the storage system.

When enabled, you agree to allow the feature to remotely monitor the storage system, collect diagnostic information, and transmit the data to a remote support server. Each time SupportAssist runs, a service tag is sent that includes a unique identifier for your system. This identifier can be used to contact you if your system needs repair.

Table 13. Capabilities by service contract type

Capability	Description	Basic Hardware	ProSupport	ProSupport Plus
Automated collection of system information	The system information that is required for troubleshooting an issue is automatically collected from the monitored device and sent securely to the backend.	✓	✓	✓
Proactive detection of hardware failures	Receives alerts for hardware events that occur in monitored devices and proactively determines if the alerts indicate a hardware failure.	✓	✓	✓
Automated support case creation	When a hardware failure is detected either proactively or predictively, a support case is automatically created with technical support.	✗	✓	✓
Automated email notification	An email notification about the support case or issue is automatically sent to the primary and secondary contacts.	✗	✓	✓
Proactive response from technical support	A technical support agent contacts you proactively about the support case and helps you resolve the issue.	✗	✓	✓
Proactive parts dispatch	After analyzing the collected system information, if the technical support agent	✗	✓	✓

Table 13. Capabilities by service contract type (continued)

Capability	Description	Basic Hardware	ProSupport	ProSupport Plus
	determines that a part must be replaced to resolve the issue, a replacement part is dispatched to you based on the dispatch preferences.			
Predictive detection of hardware failures	Intelligent analysis of system information that is collected from a monitored device is used to predict hardware failures that may occur in future.	✘	✘	✔

Table 14. Port requirements

Type of connection	Ports that must be open	Protocol used	Communication
Connect Directly	443	TCP	Outbound
Connect via Gateway Server	9443	TCP	Outbound

Enable SupportAssist

If you did not enable SupportAssist during the guided setup, you can enable it from the **Maintenance > Support** panel.

About this task

Ensure that the ME5 Series storage system meets the network requirement for a secure gateway. See [SupportAssist direct connection requirements](#).


Steps

- Go to **Maintenance > Support**.
- Select the **SupportAssist** check box.
The SupportAssist agreement is displayed.
- Read through the agreement, then acknowledge it by selecting **I accept this agreement** and clicking **ACCEPT AND CONTINUE**.
- In the **Connection Information** panel, select your connectivity options:
 - Connection Type.** Select whether to connect directly or through a gateway.
 - Proxy Details.** If applicable, select **Use a Proxy Server** and then enter the server settings.
 - Access key and PIN.** Enter the information requested. If you do not have the access key or PIN, click on **request a new Access Key and PIN** and follow the prompts to have new key information emailed to you.
- Click **Test and Enable Connectivity**.
Test results are displayed, you can either go back and re-enter information, or click **Continue** to proceed.
- In the **Contact Information** panel, enter the primary contact information and select the preferred contact settings. You can also enter information for a secondary contact
- Click **Apply changes**, and click **OK** on the confirmation prompt.
To disable SupportAssist:
 - Clear the **SupportAssist** check box.
The SupportAssist opt out confirmation panel is displayed.
 - Click **Yes**.
 - Click **Apply** or **Apply and Close**.
The SupportAssist changes confirmation panel is displayed.
 - Click **OK**.

Configure SupportAssist to use a web proxy

If the storage array does not have direct access to the Internet, SupportAssist can use a web proxy to send data to technical support .

About this task

 **NOTE:** DNS must be configured on each controller to allow SupportAssist to work with a web proxy. To configure DNS, see [Configuring DNS settings](#).

Steps

1. Perform one of the following actions to access the SupportAssist options:
 - In the Home topic, select **Action > System Settings**, then click the **SupportAssist** tab.
 - In the System topic, select **Action > System Settings**, then click the **SupportAssist** tab.
 - In the Welcome panel, select **System Settings**, then click the **SupportAssist** tab.
2. On the **Web Proxy Settings** tab:
 - a. Select the **Web Proxy** checkbox.
 - b. Type the hostname IP address of the proxy server in the **Host** field.
 - c. Type the port number of the proxy server in the **Port** field.
 - d. If the proxy server requires authentication, type the credentials in the **User Name** and **Password** fields.
3. Click **Apply** or **Apply and Close**, and click **OK** on the confirmation panel.

Add support contacts

Add support contacts from the SupportAssist panel.

Steps

1. Go to **Maintenance > Support**.
2. Expand **Support Contacts**.
3. Enter the information in the fields provided.
4. To add another contact, click **Add Secondary Contact**.
5. Click **Apply Changes**, and click **OK** on the confirmation prompt.

Set advanced support options

The **Advanced** section of the SupportAssist tab contains options for interacting with support and for putting the system into maintenance mode.

Steps

1. Go to **Maintenance > Support**.
2. Expand **Advanced**.

The following options are available:

 - **Maintenance Mode**—select to prevent support cases from being created during maintenance operations such as firmware updates.
 - **Automatic Case Creation**—select to have SupportAssist automatically create support cases when specific events occur.
 - **Re-authenticate SupportAssist**—enter your Access Key and PIN number to refresh your connection to SupportAssist. You can also request a new Access Key and PIN from the link provided.
3. Click **Apply Changes**, and click **OK** on the confirmation prompt.

Enable or disable CloudIQ


CloudIQ is enabled by default on ME5 Series storage systems. To send data to CloudIQ, you must have a business account on dell.com. In addition, the ME5 Series storage system must be onboarded to CloudIQ and SupportAssist must be enabled on the storage system.

About this task

To stop sending data to CloudIQ without removing the storage system from CloudIQ, clear the **Enable CloudIQ** checkbox.

Steps

1. Go to **Maintenance > Support**.
2. On the **SupportAssist** tab, select or clear the **Enable CloudIQ** checkbox.
3. Click **Apply Changes**, and click **OK** at the confirmation prompt.

 **NOTE:** It may take several hours for changes to the CloudIQ setting to take effect.

Other management interfaces

The ME5 Series system supports SNMP, FTP, SFTP and SLP protocols for interacting with the system.

Topics:

- [SNMP reference](#)
- [Using FTP and SFTP](#)
- [Using SLP](#)

SNMP reference

This appendix describes the Simple Network Management Protocol (SNMP) capabilities that Dell storage systems support. This includes standard MIB-II, the FibreAlliance SNMP Management Information Base (MIB) version 2.2 objects, and enterprise traps.

The storage systems can report their status through SNMP. SNMP provides basic discovery using MIB-II, more detailed status with the FA MIB 2.2, and asynchronous notification using enterprise traps.


SNMP is a widely used network monitoring and control protocol. It is an application layer protocol that facilitates the exchange of management information between network devices. It is part of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite.

SNMP enables network administrators to manage network performance, find and solve network problems, and plan for network growth. Data is passed from SNMP agents reporting activity on each network device to the workstation console used to oversee the network. The agents return information contained in a Management Information Base (MIB), which is a data structure that defines what is obtainable from the device and what can be controlled (turned on and off, etc.).

Supported SNMP versions

The storage systems allow use of SNMPv2c or SNMPv3. SNMPv2c uses a community-based security scheme. For improved security, SNMPv3 provides authentication of the network management system that is accessing the storage system, and encryption of the information transferred between the storage system and the network management system.

When SNMPv3 is disabled, SNMPv2c will be active. When SNMPv3 is enabled, SNMPv2c will only have access to the MIB-II common system information. This allows device discovery.

When SNMPv3 is disabled, SNMPv2c will be active. When SNMPv3 is enabled, SNMPv2c will be inactive. To enable SNMPv3, create a user with the snmpuser interface (**Settings > Users > SNMPv3 > Add New SNMPv3 User**). To disable SNMPv3, delete all SNMPv3 users (**Settings > Users > SNMPv3 > **).

Whether you use SNMPv2c or v3, note that the only SNMP-writable information is the system contact, name, and location. System data, configuration, and state cannot be changed via SNMP.

Standard MIB-II behavior

MIB-II is implemented to support basic discovery and status.

An SNMP object identifier (OID) is a number assigned to devices in a network for identification purposes. OID numbering is hierarchical. Using the IETF notation of digits and dots resembling very long IP addresses, various registries such as ANSI assign high-level numbers to vendors and organizations. They, in turn, append digits to the number to identify individual devices or software processes.

In the system group, all objects can be read. The contact, name, and location objects can be set.

In the interfaces group, an internal PPP interface is documented, but it is not reachable from external to the device.

The address translation (at) and external gateway protocol (egp) groups are not supported.

Enterprise traps

Traps can be generated in response to events occurring in the storage system. These events can be selected by severity and by individual event type. A maximum of three SNMP trap destinations can be configured by IP address.

Enterprise event severities are informational, minor, major, and critical. There is a different trap type for each of these severities. The trap format is represented by the enterprise traps MIB. Information included is the event ID, the event code type, and a text description generated from the internal event. Equivalent information can also be sent using email or popup alerts to users who are logged in to the PowerVault Manager.

FA MIB 2.2 SNMP behavior

The FA MIB 2.2 objects are in compliance with the FibreAlliance MIB v2.2 Specification (FA MIB2.2 Spec).

FA MIB 2.2 was never formally adopted as a standard, but it is widely implemented and contains many elements useful for storage products. This MIB generally does not reference and integrate with other standard SNMP information. It is implemented under the experimental subtree.

Significant status within the device includes such elements as its temperature and power sensors, the health of its storage elements such as virtual disks, and the failure of any redundant component including an I/O controller. While sensors can be individually queried, for the benefit of network management systems all the above elements are combined into an overall status sensor. This is available as the unit status (`connUnitStatus` for the only unit).

The revisions of the various components within the device can be requested through SNMP.

The port section is only relevant to products with Fibre Channel host ports.

The event table allows 400 recently-generated events to be requested. Informational, minor, major, or critical event types can be selected. Whichever type is selected enables the capture of that type and more severe events. This mechanism is independent of the assignment of events to be generated into traps.

The traps section is not supported. It has been replaced by an ability to configure trap destinations using the CLI or the PowerVault Manager. The statistics section is not implemented.

The following table lists the MIB objects, their descriptions and the value set in ME5 Series storage systems. Unless specified otherwise, objects cannot be set.

Table 15. FA MIB 2.2 objects, descriptions, and values

Object	Description	Value
<code>RevisionNumber</code>	Revision number for this MIB	0220
<code>UNumber</code>	Number of connectivity units present	1
<code>SystemURL</code>	Top-level URL of this device, for example, <code>http://10.1.2.3</code> . If a web server is not present on the device, this string is empty in accordance with the FA MIB2.2 Spec.	Default: <code>http://10.0.0.1</code>
<code>StatusChangeTime</code>	<code>sysuptime</code> timestamp of the last status change event, in centiseconds. <code>sysuptime</code> starts at 0 when the Storage Controller boots and keeps track of the up time. <code>statusChangeTime</code> is updated each time an event occurs.	0 at startup
<code>ConfigurationChangeTime</code>	<code>sysuptime</code> timestamp of the last configuration change event, in centiseconds. <code>sysuptime</code> starts at 0 when the Storage Controller boots and keeps track of the up time. <code>configurationChangeTime</code> is updated each time an event occurs	0 at startup
<code>ConnUnitTableChangeTime</code>	<code>sysuptime</code> timestamp of the last update to the <code>connUnitTable</code> (an	0 always (entries are not added to or deleted from the <code>connUnitTable</code>)

Table 15. FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
	entry was either added or deleted),in centiseconds	
connUnitTable	Includes the following objects as specified by the FA MIB2.2 Spec:	
connUnitId	Unique identification for this connectivity unit	Total of 16 bytes comprised of 8 bytes of the node WWN or similar serial number-based identifier (for example,1000005013b05211) with the trailing 8 bytes equal to zero
connUnitGlobalId	Same as connUnitId	Same as connUnitId
connUnitType	Type of connectivity unit	storage-subsystem (11)
connUnitNumports	Number of host ports in the connectivity unit	Number of host ports
connUnitState	Overall state of the connectivity unit	online (2) or unknown (1), as appropriate
connUnitStatus	Overall status of the connectivity unit	ok (3), warning (4), failed (5), or unknown (1), as appropriate
connUnitProduct	Connectivity unit vendor product model name	Model string
connUnitSn	Serial number for this connectivity unit	Serial number string
connUnitUpTime	Number of centiseconds since the last unit initialization	0 at startup
connUnitUrl	Same as systemURL	Same as systemURL
connUnitDomainId	Not used; set to all 1s as specified by the FA MIB2.2 Spec	0xFFFF
connUnitProxyMaster	Stand-alone unit returns yes for this object	yes (3) since this is a stand-alone unit
connUnitPrincipal	Whether this connectivity unit is the principal unit within the group of fabric elements. If this value is not applicable, returns unknown.	unknown (1)
connUnitNumSensors	Number of sensors in the connUnitSensorTable	33
connUnitStatusChangeTime	Same as statusChangeTime	Same as statusChangeTime
connUnitNumRevs	Number of revisions in the connUnitRevsTable	16
connUnitNumZones	Not supported	0
connUnitModuleId	Not supported	16 bytes of 0s
connUnitName	Settable: Display string containing a name for this connectivity unit	Default: Uninitialized Name
connUnitInfo	Settable: Display string containing information about this connectivity unit	Default: Uninitialized Info
connUnitControl	Not supported	invalid (2) for an SNMP GET operation and not settable through an SNMP SET operation.
connUnitContact	Settable: Contact information for this connectivity unit	Default: Uninitialized Contact

Table 15. FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
connUnitLocation	Settable: Location information for this connectivity unit	Default: Uninitialized Location
connUnitEventFilter	Defines the event severity that will be logged by this connectivity unit. Settable only through the PowerVault Manager.	Default: info (8)
connUnitNumEvents	Number of events currently in the connUnitEventTable	Varies as the size of the Event Table varies
connUnitMaxEvents	Maximum number of events that can be defined in the connUnitEventTable	400
connUnitEventCurrID	Not supported	0
connUnitRevsTable	Includes the following objects as specified by the FA MIB2.2 Spec:	
connUnitRevsUnitId	connUnitId of the connectivity unit that contains this revision table	Same as connUnitId
connUnitRevsIndex	Unique value for each connUnitRevsEntry between 1 and connUnitNumRevs	See connUnitRevsDescription Table
connUnitRevsRevId	Vendor-specific string identifying a revision of a component of the connUnit	String specifying the code version. Reports "Not Installed or Offline" if module information is not available.
connUnitRevsDescription	Display string containing description of a component to which the revision corresponds	See connUnitRevsDescription Table
connUnitSensorTable	Includes the following objects as specified by the FA MIB2.2 Spec:	
connUnitSensorUnitId	connUnitId of the connectivity unit that contains this sensor table	Same as connUnitId
connUnitSensorIndex	Unique value for each connUnitSensorEntry between 1 and connUnitNumSensors	See External details for connUnitSensor Table
connUnitSensorName	Display string containing textual identification of the sensor intended primarily for operator use	See External details for connUnitSensor Table
connUnitSensorStatus	Status indicated by the sensor	ok (3), warning (4), or failed (5) as appropriate for FRUs that are present, or other (2) if FRU is not present.
connUnitSensorInfo	Not supported	Empty string
connUnitSensorMessage	Description the sensor status as a message	connUnitSensorName followed by the appropriate sensor reading. Temperatures display in both Celsius and Fahrenheit. For example, CPU Temperature (Controller Module A): 48C 118F). Reports "Not installed" or "Offline" if data is not available.
connUnitSensorType	Type of component being monitored by this sensor	See External details for connUnitSensor Table
connUnitSensorCharacteristic	Characteristics being monitored by this sensor	See External details for connUnitSensor Table
connUnitPortTable	Includes the following objects as specified by the FA MIB2.2 Spec:	

Table 15. FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
connUnitPortUnitId	connUnitId of the connectivity unit that contains this port	Same as connUnitId
connUnitPortIndex	Unique value for each connUnitPortEntry between 1 and connUnitNumPorts	Unique value for each port, between 1 and the number of ports
connUnitPortType	Port type	not-present (3), or n-port (5) for point-to-point topology, or l-port (6)
connUnitPortFCClassCap	Bit mask that specifies the classes of service capability of this port. If this is not applicable, returns all bits set to zero.	Fibre Channel ports return 8 for class-three
connUnitPortFCClassOp	Bit mask that specifies the classes of service that are currently operational. If this is not applicable, returns all bits set to zero.	Fibre Channel ports return 8 for class-three
connUnitPortState	State of the port hardware	unknown (1), online (2), offline (3), bypassed (4)
connUnitPortStatus	Overall protocol status for the port	unknown (1), unused (2), ok (3), warning (4), failure (5), not participating (6), initializing (7), bypass (8)
connUnitPortTransmitterType	Technology of the port transceiver	unknown (1) for Fibre Channel ports
connUnitPortModuleType	Module type of the port connector	unknown (1)
connUnitPortWwn	Fibre Channel World Wide Name (WWN) of the port if applicable	WWN octet for the port, or empty string if the port is not present
connUnitPortFCId	Assigned Fibre Channel ID of this port	Fibre Channel ID of the port. All bits set to 1 if the Fibre Channel ID is not assigned or if the port is not present
connUnitPortSn	Serial number of the unit (for example, for a GBIC). If this is not applicable, returns an empty string.	Empty string
connUnitPortRevision	Port revision (for example, for a GBIC)	Empty string
connUnitPortVendor	Port vendor (for example, for a GBIC)	Empty string
connUnitPortSpeed	Speed of the port in KByte per second (1 KByte = 1000 Byte)	Port speed in KByte per second, or 0 if the port is not present
connUnitPortControl	Not supported	invalid (2) for an SNMP GET operation and not settable through an SNMP SET operation
connUnitPortName	String describing the addressed port	See External details for connUnitPortTable
connUnitPortPhysicalNumber	Port number represented on the hardware	Port number represented on the hardware
connUnitPortStatObject	Not supported	0 (No statistics available)
connUnitEventTable	Includes the following objects as specified by the FA MIB2.2 Spec:	
connUnitEventUnitId	connUnitId of the connectivity unit that contains this port	Same as connUnitId

Table 15. FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
connUnitEventIndex	Index into the event buffer of the connectivity unit, incremented for each event	Starts at 1 every time there is a table reset or the unit event table reaches its maximum index value
connUnitEventId	Internal event ID, incremented for each event, ranging between 0 and connUnitMaxEvents	Starts at 0 every time there is a table reset or connUnitMaxEvents is reached
connUnitREventTime	Real time when the event occurred, in the following format: DDMMYYYY HHMMSS	0 for logged events that occurred prior to or at startup
connUnitSEventTime	sysuptime timestamp when the event occurred	0 at startup
connUnitEventSeverity	Event severity level	error (5), warning (6) or info (8)
connUnitEventType	Type of this event	As defined in CAPI
connUnitEventObject	Not used	0
connUnitEventDescr	Text description of this event	Formatted event, including relevant parameters or values
connUnitLinkTable	Not supported	N/A
connUnitPortStatFabricTable	Not supported	N/A
connUnitPortStatSCSITable	Not supported	N/A
connUnitPortStatLANTable	Not supported	N/A
SNMP Traps	The following SNMP traps are supported:	
trapMaxClients	Maximum number of trap clients	1
trapClientCount	Number of trap clients currently enabled	1 if traps enabled; 0 if traps not enabled
connUnitEventTrap	This trap is generated each time an event occurs that passes the connUnitEventFilter and the trapRegFilter	N/A
trapRegTable	Includes the following objects per the FA MIB2.2 Spec:	
trapRegIpAddress	IP address of a client registered for traps	IP address set by user
trapRegPort	User Datagram Protocol (UDP) port to send traps to for this host	162
trapRegFilter	Settable: Defines the trap severity filter for this trap host. The connUnit will send traps to this host that have a severity level less than or equal to this value.	Default: warning (6)
trapRegRowState	Specifies the state of the row	READ: rowActive (3) if traps are enabled. Otherwise rowInactive(2) WRITE: Not supported

External details for certain FA MIB 2.2 objects

Tables in this section specify values for certain objects described in the previous table.

Table 16. connUnitRevsTable index and description values

connUnitRevsIndex	connUnitRevsDescription
1	CPU Type for Storage Controller (Controller A)
2	Bundle revision for Controller (Controller A)
3	Build date for Storage Controller (Controller A)
4	Code revision for Storage Controller (Controller A)
5	Code base level for Storage Controller (Controller A)
6	FPGA code revision for Memory Controller (Controller A)
7	Loader code revision for Storage Controller (Controller A)
8	CAPI revision (Controller A)
9	Code revision for Management Controller (Controller A)
10	Loader code revision for Management Controller (Controller A)
11	Code revision for Expander Controller (Controller A)
12	CPLD code revision (Controller A)
13	Hardware revision (Controller A)
14	Host interface module revision (Controller A)
15	HIM revision (Controller A)
16	Backplane type (Controller A)
17	Host interface hardware (chip) revision (Controller A)
18	Disk interface hardware (chip) revision (Controller A)
19	CPU Type for Storage Controller (Controller B)
20	Bundle revision for Controller (Controller B)
21	Build date for Storage Controller (Controller B)
22	Code revision for Storage Controller (Controller B)
23	Code base level for Storage Controller (Controller B)
24	FPGA code revision for Memory Controller (Controller B)
25	Loader code revision for Storage Controller (Controller B)
26	CAPI revision (Controller B)
27	Code revision for Management Controller (Controller B)
28	Loader code revision for Management Controller (Controller B)
29	Code revision for Expander Controller (Controller B)
30	CPLD code revision (Controller B)
31	Hardware revision (Controller B)
32	Host interface module revision (Controller B)
33	HIM revision (Controller B)
34	Backplane type (Controller B)

Table 16. connUnitRevsTable index and description values (continued)

connUnitRevsIndex	connUnitRevsDescription
35	Host interface hardware (chip) revision (Controller B)
36	Disk interface hardware (chip) revision (Controller B)

External details for connUnitSensorTable

Table 17. connUnitSensorTable index, name, type, and characteristic values

connUnitSensorIndex	connUnitSensorName	connUnitSensorType	connUnitSensorCharacteristic
1	Onboard Temperature 1 (Controller A)	board(8)	temperature(3)
2	Onboard Temperature 1 (Controller B)	board(8)	temperature(3)
3	Onboard Temperature 2 (Controller A)	board(8)	temperature(3)
4	Onboard Temperature 2 (Controller B)	board(8)	temperature(3)
5	Onboard Temperature 3 (Controller A)	board(8)	temperature(3)
6	Onboard Temperature 3 (Controller B)	board(8)	temperature(3)
7	Disk Controller Temperature (Controller A)	board(8)	temperature(3)
8	Disk Controller Temperature (Controller B)	board(8)	temperature(3)
9	Memory Controller Temperature (Controller A)	board(8)	temperature(3)
10	Memory Controller Temperature (Controller B)	board(8)	temperature(3)
11	Capacitor Pack Voltage (Controller A)	board(8)	power(9)
12	Capacitor Pack Voltage (Controller B)	board(8)	power(9)
13	Capacitor Cell 1 Voltage (Controller A)	board(8)	power(9)
14	Capacitor Cell 1 Voltage (Controller B)	board(8)	power(9)
15	Capacitor Cell 2 Voltage (Controller A)	board(8)	power(9)
16	Capacitor Cell 2 Voltage (Controller B)	board(8)	power(9)
17	Capacitor Cell 3 Voltage (Controller A)	board(8)	power(9)
18	Capacitor Cell 3 Voltage (Controller B)	board(8)	power(9)
19	Capacitor Cell 4 Voltage (Controller A)	board(8)	power(9)
20	Capacitor Cell 4 Voltage (Controller B)	board(8)	power(9)
21	Capacitor Charge Percent (Controller A)	board(8)	other(2)
22	Capacitor Charge Percent (Controller B)	board(8)	other(2)
23	Overall Status	enclosure(7)	other(2)
24	Upper IOM Temperature (Controller A)	enclosure(7)	temperature(3)
25	Lower IOM Temperature (Controller B)	enclosure(7)	temperature(3)
26	Power Supply 1 (Left) Temperature	power-supply(5)	temperature(3)
27	Power Supply 2 (Right) Temperature	power-supply(5)	temperature(3)

Table 17. connUnitSensorTable index, name, type, and characteristic values (continued)

connUnitSensorIndex	connUnitSensorName	connUnitSensorType	connUnitSensorCharacteristic
28	Upper IOM Voltage, 12V (Controller A)	enclosure(7)	power(9)
29	Upper IOM Voltage, 5V (Controller A)	enclosure(7)	power(9)
30	Lower IOM Voltage, 12V (Controller B)	enclosure(7)	power(9)
31	Lower IOM Voltage, 5V (Controller B)	enclosure(7)	power(9)
32	Power Supply 1 (Left) Voltage, 12V	power-supply(5)	power(9)
33	Power Supply 1 (Left) Voltage, 5V	power-supply(5)	power(9)
34	Power Supply 1 (Left) Voltage, 3.3V	power-supply(5)	power(9)
35	Power Supply 2 (Right) Voltage, 12V	power-supply(5)	power(9)
36	Power Supply 2 (Right) Voltage, 5V	power-supply(5)	power(9)
37	Power Supply 2 (Right) Voltage, 3.3V	power-supply(5)	power(9)
38	Upper IOM Voltage, 12V (Controller A)	enclosure(7)	currentValue(6)
39	Lower IOM Voltage, 12V (Controller B)	enclosure(7)	currentValue(6)
40	Power Supply 1 (Left)Current,12V	power-supply(5)	currentValue(6)
41	Power Supply 1 (Left)Current, 5V	power-supply(5)	currentValue(6)
42	Power Supply 2 (Right)Current,12V	power-supply(5)	currentValue(6)
43	Power Supply 2 (Right)Current, 5V	power-supply(5)	currentValue(6)

External details for connUnitPortTable

Table 18. connUnitPortTable index and name values

connUnitPortIndex	connUnitPortName
0	Host Port 0 (Controller A)
1	Host Port 1 (Controller B)
2	Host Port 2 (Controller B)
3	Host Port 3 (Controller B)

Configuring SNMP event notification in the PowerVault Manager

Steps

1. Verify that SNMP service is enabled on the storage system. See [Enabling or disabling system management services](#).
2. Configure and enable SNMP traps. See [Managing SNMPv3 users](#).
3. Optionally, configure a user account to receive SNMP traps. See [Managing SNMPv3 users](#).

SNMP management

You can manage storage devices using SNMP with a network management system such as HPE Systems Insight Manager (SIM) or HP Instant Support Enterprise Edition (ISEE). See their documentation for information about loading MIBs, configuring events, and viewing and setting group objects.

To view and set system group objects, SNMP must be enabled in the storage system. See [Enabling or disabling system management services](#). To use SNMPv3, it must be configured in both the storage system and the network management system that intends to access the storage system or receive traps from it. In the storage system, SNMPv3 is configured through the creation and use of SNMP user accounts, as described in [User settings](#). The same users, security protocols, and passwords must be configured in the network management system.

To obtain the MIB, see www.dell.com/support.

Using FTP and SFTP

Although the PowerVault Manager is the preferred interface for downloading log data and historical disk-performance statistics, you can also use FTP and SFTP to do these tasks, and to install security certificates and keys. FTP and SFTP can also be used for updating firmware and installing security certificates and keys.

NOTE: Dell recommends using SFTP rather than FTP because SFTP is a secured protocol.

NOTE: Do not attempt to do more than one of the operations in this appendix at the same time. They can interfere with each other and the operations may fail. Specifically, do not try to do more than one firmware update at the same time or try to download system logs while doing a firmware update.

Downloading system logs

To help service personnel diagnose a system problem, you might be asked to provide system log data.

You can download this data by accessing the system's FTP or SFTP interface and running the `get logs` command. When both controllers are online, regardless of operating mode, `get logs` will download a single, compressed zip file that includes:

- Device status summary, which includes basic status and configuration data for the system
- MC logs from each controller
- Event log from each controller
- Debug log from each controller
- Boot log from each controller, which shows the startup sequence
- Critical error dumps from each controller, if critical errors have occurred
- CAPI traces from each controller

Use a command-line-based FTP/SFTP client. A UI-based FTP/SFTP client might not work.

Download system logs

Perform the following steps to download the system logs:

Steps

1. In the PowerVault Manager, prepare to use FTP/SFTP:
 - a. Determine the network-port IP addresses of the system controllers. See [Configuring controller network ports](#).
 - b. Verify that the FTP/SFTP service is enabled on the system. See [Enabling or disabling system-management settings](#).
 - c. Verify that the user you will log in as has permission to use the FTP interface. The same setting allows a user to transfer files using both FTP and SFTP. See [User settings](#).
2. Open a Command Prompt (Windows) or a terminal window (UNIX) and go to the destination directory for the log file.
3. Using the FTP/SFTP port specified in the system services settings, enter:

```
sftp -P <port> <controller-network-address> or  
ftp <controller-network-address>
```

For example:

```
sftp -P 1022 10.235.216.152
```

or

```
ftp 10.1.0.9
```

4. Log in as a user that has permission to use the FTP/SFTP interface.

5. Make sure the client is in binary transfer mode. Enter:

```
binary
```

6. Enter:


```
get logs <filename>.zip
```

where <filename> is the file that contains the logs. Dell recommends using a filename that identifies the system, controller, and date.

```
get logs Storage2_A_20120126.zip
```

In FTP, wait for the message `Operation Complete` to appear. No messages are displayed in SFTP; instead, the `get` command returns once the logs collection is finished.

7. Quit the FTP/SFTP session.

 **NOTE:** The log files must be extracted from .zip file to view them. To examine diagnostic data, view `store_yyyy_mm_dd__hh_mm_ss.logs` file first.

Transferring log data to a log-collection system

If the log-management feature is configured in pull mode, a log-collection system can access the FTP or SFTP interface on the storage system and use the `get managed-logs` command to retrieve untransferred data from a system log file.

This command retrieves the untransferred data from the specified log to a compressed zip file on the log-collection system. Following the transfer of the log data, the log capacity status is reset to zero indicate that there is no untransferred data. Log data is controller specific.

For an overview of the log-management feature, see [Managed logs](#).

Use a command-line-based FTP/SFTP client. A UI-based FTP client might not work.

Transfer log data to a log-collection system

Perform the following steps to transfer log data to a log-collection system:

Steps

1. In the PowerVault Manager, prepare to use FTP/SFTP:

- Determine the network-port IP addresses of the system controllers. See [Configuring controller network ports](#).
- Verify that the FTP/SFTP service is enabled on the system. See [Enabling or disabling system-management settings](#).
- Verify that the user you will log in as has permission to use the FTP interface. See [User settings](#).

2. On the log-collection system, open a Command Prompt (Windows) or a terminal window (UNIX) and go to the destination directory for the log file.

3. Enter:

```
sftp -P <port> <controller-network-address> or  
ftp <controller-network-address>
```

For example:

```
sftp -P 1022 10.235.216.152
```

or

```
ftp 10.1.0.9
```

4. Log in as a user that has permission to use the FTP/SFTP interface.

5. Make sure the client is in binary transfer mode. Enter:

```
binary
```

6. Enter:


```
get managed-logs:<log-type> <filename>.zip where:
```

- `log-type` specifies the type of log data to transfer:
 - `crash1`, `crash2`, `crash3`, or `crash4`: One of the four crash logs on the storage controller.
 - `ecdebug`: Expander Controller log.
 - `mc`: Management Controller log.
 - `scdebug`: Storage Controller log.
- `<filename>` is the file that contains the transferred data. Dell recommends using a filename that identifies the system, controller, and date.

```
get managed-logs:scdebug Storage2-A_scdebug_2011_08_22.zip
```

In FTP, wait for the message `operation Complete` to appear. No messages are displayed in SFTP; instead, the `get` command returns once the data transfer is finished.

7. Quit the FTP/SFTP session.

 **NOTE:** The log files must be extracted from .zip file to view them.

Downloading historical disk-performance statistics

You can access the FTP or SFTP interface on the storage system and use the `get perf` command to download historical disk-performance statistics for all disks in the storage system. This command downloads the data in CSV format to a file, for import into a spreadsheet or other third-party application.

The number of data samples downloaded is fixed at 100 to limit the size of the data file to be generated and transferred. The default is to retrieve all the available data (up to six months) aggregated into 100 samples. You can specify a different time range by specifying a start and end time. If the specified time range spans more than 100 15-minute samples, the data will be aggregated into 100 samples.

The resulting file will contain a row of property names and a row for each data sample, as shown in the following example. For property descriptions, see the topic about the `disk-hist-statistics` basetype in the CLI Reference Guide.

```
"sample-time","durable-id","serial-number","number-of-ios", ...
"2012-01-26 01:00:00","disk_1.1","PLV2W1XE","2467917", ...
"2012-01-26 01:15:00","disk_1.1","PLV2W1XE","2360042", ...
...
```

Use a command-line-based FTP/SFTP client. A UI-based FTP/SFTP client might not work.

Retrieve historical disk-performance statistics

Perform the following steps to retrieve historical disk-performance statistics:

Steps

1. In the PowerVault Manager, prepare to use FTP/SFTP:
 - a. Determine the network-port IP addresses of the system controllers. See [Configuring controller network ports](#).
 - b. Verify that the FTP/SFTP service is enabled on the system. See [Enabling or disabling system-management settings](#).
 - c. Verify that the user you will log in as has permission to use the FTP interface. See [User settings](#).
2. Open a Command Prompt (Windows) or a terminal window (UNIX) and go to the destination directory for the log file.
3. Type:

```
sftp -P <port> <controller-network-address> or
ftp <controller-network-address>
```

For example:

```
sftp -P 1022 10.235.216.152 or
ftp 10.1.0.9
```

4. Make sure the client is in binary transfer mode. Enter:

```
binary
```

5. Log in as a user that has permission to use the FTP/SFTP interface.

6. Enter:

```
get perf:<date/time-range> <filename.csv> where:
```

- *date/time-range* is optional and specifies the time range of data to transfer, in the format: *start.yyyy-mm-dd.hh:mm.[AM|PM].end.yyyy-mm-dd.hh:mm.[AM|PM]*. The string must contain no spaces.
- *filename.csv* is the file that contains the data. Dell recommends using a filename that identifies the system, controller, and date.

```
get perf:start.2019-01-26.12:00.PM.end.2019-01-26.23:00.PM Storage2_A_20120126.csv
```

In FTP, wait for the message `Operation Complete` to appear. No messages are displayed in SFTP; instead, the `get` command returns once the download is finished.

7. Quit the FTP/SFTP session.

Updating firmware

As a user in the `manage` role, you can update the versions of firmware in controller modules, expansion modules in drive enclosures, and disks.

NOTE: To ensure success of an online update, select a period of low I/O activity. This helps the update complete as quickly as possible and avoids disruptions to host and applications due to timeouts. Attempting to update a storage system that is processing a large, I/O-intensive batch job will likely cause hosts to lose connectivity with the storage system.

NOTE: Consider the following points before performing a firmware update:

- If a disk group is quarantined, resolve the problem that is causing the component to be quarantined before updating the firmware.
- If any unwritten cache data is present, firmware update will not proceed. Before you can update firmware, unwritten data must be removed from cache. For more information about the `clear cache` command, see the CLI Reference Guide.
- If the health of the system is `Fault`, the firmware update cannot proceed. Before you can update firmware, you must resolve the problem specified by the Health Reason value on the System Overview panel.

Updating controller-module firmware

In a dual-controller system, both controllers should run the same firmware version. Storage systems in a replication set should run the same or compatible firmware versions. You can update the firmware in each controller module by loading a firmware file obtained from the enclosure vendor.

If you have a dual-controller system and the Partner Firmware Update (PFU) option is enabled, when you update one controller the system automatically updates the partner controller. If PFU is disabled, after updating firmware on one controller you must log into the IP address of the partner controller and perform this firmware update on that controller also.

For best results, ensure the storage system is in a healthy state before starting firmware update.

NOTE: For information about supported releases for firmware update, see the Release Notes.

Update controller module firmware

Perform the following steps to update the controller module firmware:

Steps

1. Obtain the appropriate firmware file, and download it to your computer or network.
2. In the PowerVault Manager, prepare to use FTP/SFTP:
 - a. Determine the network-port IP addresses of the system controllers.
 - b. Verify that the FTP/SFTP service is enabled on the system.
 - c. Verify that the user you will log in as has permission to use the FTP/SFTP interface. The same setting allows a user to transfer files using both FTP and SFTP.

3. Open a Command Prompt (Windows) or a terminal window (UNIX) and go to the directory containing the firmware file to load.

4. Enter:

```
sftp -P <port> <controller-network-address> or  
ftp <controller-network-address>
```

For example:

```
sftp -P 1022 10.235.216.152 or  
ftp 10.1.0.9
```

5. Log in as a user with manage role permissions and FTP/SFTP interface permissions.

6. Make sure that the client is in binary transfer mode. Enter:

```
binary
```

7. Enter:

```
put <firmware-file> flash
```

CAUTION: Do not perform a power cycle or controller restart during a firmware update. If the update is interrupted or there is a power failure, the module might become inoperative. If this issue occurs, contact technical support. The module might need to be returned to the factory for reprogramming.

NOTE: If you attempt to load an incompatible firmware version, the message `*** Code Load Fail. Bad format image. ***` is displayed and after a few seconds the FTP/SFTP prompt is redisplayed. The code is not loaded.

Firmware update typically takes 10 minutes for a controller having current CPLD firmware, or 20 minutes for a controller with downlevel CPLD firmware. If the controller enclosure has attached enclosures, allow additional time for each enclosure management processor (EMP) to be updated in the expansion module. It typically takes 2.5 minutes for each EMP in a drive enclosure.

NOTE: If you are using a Windows FTP/SFTP client, during firmware update a client-side FTP/SFTP application issue or timeout setting can cause the FTP/SFTP session to be aborted. If this issue persists, try using the PowerVault Manager to perform the update, use another client, or use another FTP/SFTP application.

If the Storage Controller cannot be updated, the update operation is canceled. If the FTP/SFTP prompt does not return, quit the FTP/SFTP session and log in again. Verify that you specified the correct firmware file and repeat the update. If this problem persists, contact technical support.

When firmware update on the local controller is complete, the FTP session returns to the `sftp>` prompt, and the FTP/SFTP session to the local MC is closed. Use a management interface to monitor the system and determine when the update is complete.

If the Partner Firmware Update (PFU) feature is enabled, both controllers will update. If the PowerVault Manager is open, it will display a pop-up showing update progress. Progress can also be monitored using the `show firmware-update-status` CLI command. For more information on this command, see the CLI Reference Guide.

8. Quit the FTP/SFTP session.

9. Clear your web browser cache, and then sign in to the PowerVault Manager.

If PFU is running on the controller you sign in to, a dialog box shows PFU progress and prevents you from performing other tasks until PFU is complete.

NOTE: If PFU is enabled for the system, after firmware update has completed on both controllers, check the system health. If the system health is Degraded and the health reason indicates that the firmware version is incorrect, verify that you specified the correct firmware file. If this problem persists, contact technical support.

Updating expansion-module and drawer firmware

ME5 Series supports two expansion modules. Each expansion module contains an enclosure management processor (EMP). In an enclosure with drawers, each drawer contains two EMPs, which are also referred to as "modules." All modules of the same product model should run the same firmware version.

Expansion-module and drawer firmware is updated in either of two ways:

- When you update controller-module firmware, all expansion-module and drawer EMPs are automatically updated to a compatible firmware version.

- You can update the firmware in each expansion-module and drawer EMP by loading a firmware file obtained from the enclosure vendor.

You can specify to update all expansion modules or only specific expansion modules. If you specify to update all expansion modules and the system contains more than one type of enclosure, the update will be attempted on all enclosures in the system. The update will only succeed for enclosures whose type matches the file, and will fail for enclosures of other types.

Update expansion-module and drawer firmware

Perform the following steps to update the expansion-module and drawer firmware:

Steps

1. As a user with a `manage` role, obtain the appropriate firmware file, and download it to your computer or network.
2. If you want to update all expansion modules, go to the next step. Otherwise, in the PowerVault Manager, determine the address of each expansion module to update:
 - a. In the **Maintenance > Hardware** panel, select a drive enclosure.
 - b. In the enclosure device details pane, note each EMP bus ID and target ID values. For example, 0 and 63, and 1 and 63. Bus 0 is the bus that is native to a given controller, while bus 1 is an alternate path through the partner controller. Dell recommends performing update tasks consistently through one controller to avoid confusion.
3. In the PowerVault Manager (**Settings > Network > Services** panel), prepare to use FTP:
 - a. Determine the network-port IP addresses of the system's controllers.
 - b. Verify that the FTP service is enabled on the system.
 - c. Verify that the user you will log in as has permission to use the FTP interface. If the system has a single controller, stop I/O to disk groups before starting the firmware update.
4. Open a Command Prompt (Windows) or a terminal window (UNIX) and go to the directory containing the firmware file to load.

5. Enter:

```
sftp -P<port controller-network-address> or ftp <controller-network-address>
```

For example:

```
sftp -P 1022 10.235.216.152 or
```

```
ftp 10.1.0.9
```

6. Log in as an FTP user.
7. Make sure the client is in binary transfer mode. Enter:

```
binary
```

8. Perform either of the following tasks:

- To update all expansion modules, type:

```
put <firmware-file> encl
```

- To update specific expansion modules, type:

```
put <firmware-file> encl:<EMP-bus-ID>:<EMP-target-ID>
```

CAUTION: Do not perform a power cycle or controller restart during the firmware update. If the update is interrupted or there is a power failure, the module might become inoperative. If this issue occurs, contact technical support. The module might need to be returned to the factory for reprogramming.

It typically takes 2.5 minutes to update each EMP in a drive enclosure. Wait for a message that the code load has completed.


NOTE: If the update fails, verify that you specified the correct firmware file and try the update a second time. If it fails again, contact technical support.

9. If you are updating specific expansion modules, repeat step 8 for each remaining expansion module that needs to be updated.
10. Quit the FTP session.
11. Verify that each updated expansion module has the correct firmware version.

Updating disk firmware

You can update disk firmware by loading a firmware file obtained from your reseller.

A dual-ported disk can be updated from either controller.

 **NOTE:** Disks of the same model in the storage system must have the same firmware revision.

You can specify to update all disks or only specific disks. If you specify to update all disks and the system contains more than one type of disk, the update will be attempted on all disks in the system. The update will only succeed for disks whose type matches the file, and will fail for disks of other types.

Update disk firmware

Perform the following steps to update disk firmware:

Steps

1. Obtain the appropriate firmware file, and download it to your computer or network.
2. Check the documentation from the disk manufacture to determine whether disks must be power cycled after firmware update.
3. If you want to update all disks of the type that the firmware applies to, go to the next step. Otherwise, in the PowerVault Manager, for each disk to update:
 - a. Determine the enclosure number and slot number of the disk
 - b. If the disk is associated with a disk group and is single ported, determine which controller owns the disk group.
4. In the PowerVault Manager, prepare to use FTP/SFTP:
 - a. Determine the network-port IP addresses of the system controllers.
 - b. Verify that the FTP/SFTP service is enabled on the system.
 - c. Verify that the user you plan to use has FTP interface permissions. The same setting allows a user to transfer files using both FT and SFTP.
5. Stop I/O to the storage system. During the update all volumes will be temporarily inaccessible to hosts. If I/O is not stopped, mapped hosts will report I/O errors. Volume access is restored after the update completes.
6. As a user with the `manage` role, open a Command Prompt (Windows) or a terminal window (UNIX) and go to the directory containing the firmware file to load.
7. Enter:

```
sftp -P <port> <controller-network-address> or  
ftp <controller-network-address>
```

For example:

```
sftp -P 1022 10.235.216.152 or  
ftp 10.1.0.9
```

8. Log in as an FTP/SFTP user.
9. Make sure that the client is in binary transfer mode. Enter:
`binary`
10. Perform either of the following tasks:
 - To update all disks of the type that the firmware applies to, enter:

```
put <firmware-file> disk
```

For example:

```
put AS10.bin disk
```

- To update specific disks, type:

```
put <firmware-file> disk:<enclosure-ID>:<slot-number>
```

For example:

```
put AS10.bin disk:1:11
```

CAUTION: Do not power cycle enclosures, or restart a controller during the firmware update. If the update is interrupted or there is a power failure, the disk might become inoperative. If this issue occurs, contact technical support.

It typically takes several minutes for the firmware to load. In FTP, wait for the message `Operation Complete` to appear. No messages are displayed in SFTP.

NOTE: If the update fails, verify that you specified the correct firmware file and try the update a second time. If it fails again, contact technical support.

11. If you are updating specific disks, repeat the previous step for each remaining disk to update.
12. Quit the FTP/SFTP session.
13. If the updated disks must be power cycled:
 - a. Shut down both controllers by using the PowerVault Manager.
 - b. Power cycle all enclosures as described in the *Dell PowerVault ME5 Series Storage System Deployment Guide*.
14. Verify that each disk has the correct firmware revision.

Installing a security certificate

The storage system supports use of unique certificates for secure data communications, to authenticate that the expected storage systems are being managed. Use of authentication certificates applies to the HTTPS protocol, which is used by the web server in each controller module.

As an alternative to using the CLI to create a security certificate on the storage system, you can use FTP/SFTP to install a custom certificate on the system. A certificate consists of a certificate file and an associated key file. The certificate can be created by using OpenSSL, for example, and is expected to be valid. If you replace the controller module in which a custom certificate is installed, the partner controller will automatically install the certificate file to the replacement controller module.

NOTE: Details pertaining to certificates:

- By default, the system generates a unique SSL certificate for each controller. When using FTP/SFTP to install certificates, you must explicitly install the certificate to the controller for which the session applies. For example, you cannot install a certificate to controller B during an FTP/SFTP session with controller A. You cannot install a certificate to both controllers from a given session: the installation tasks are discrete.
- The supported file format for certificates is PEM (Privacy-Enhanced Mail) encoded only.
- Supported certificate versions are x.509 v1 and v3. Avoid using unsupported versions of certificates.
- Extensions are supported if uploading certificates through FTP/SFTP. However, extensions are not supported if adding certificates using the create certificate CLI command.
- Variable key lengths are supported if uploading certificates through FTP/SFTP. However, a fixed key length of 2048 bits applies if adding certificates using the create certificate CLI command.

Two uploader roles are supported:

- The `usr` role is the default role for use by the client.
- The `mfg` role is reserved for use in engineering mode only.

Install a security certificate

Perform the following steps to install a security certificate:

Steps

1. In the PowerVault Manager, prepare to use FTP/SFTP:
 - a. Determine the network-port IP addresses of the system controllers. See [Configuring controller network ports](#).
 - b. Verify that the FTP/SFTP service is enabled on the system. See [Enabling or disabling system-management settings](#).
 - c. Verify that the user you will log in as has permission to use the FTP interface. See [Managing local users](#).
2. Open a Command Prompt (Windows) or a terminal window (UNIX) and go to the directory that contains the certificate files.
3. Enter:

```
sftp -P port controller-network-address or
```

```
ftp controller-network-address
```

For example:

```
sftp -P 1022 10.235.216.152 or  
ftp 10.1.0.9
```

4. Log in as a user with manage role permissions and FTP/SFTP interface permissions.
5. Make sure the client is in binary transfer mode. Enter:
binary
6. Enter:
put *certificate-file-name* cert-file:{usr |mfg}
where *certificate-file-name* is the name of the certificate file for your specific system.
7. Enter:
put *key-file-name* cert-key-file:{usr |mfg}
where *key-file-name* is the name of the security key file for your specific system.
8. The new security certificate takes effect.

Download a system heat map data

If requested by support engineers for analysis, you can download cumulative I/O density data, also known as heat map data, from the system. To gather this data, access the storage system FTP/SFTP interface and use the `get logs` command with the `heatmap` option to download a log file in CSV format. The file contains data for the past seven days from both controllers.

Steps

1. In the PowerVault Manager, prepare to use FTP/SFTP:
 - a. Determine the network-port IP addresses of the system controllers. See [Configuring controller network ports](#).
 - b. Verify that the FTP/SFTP service is enabled on the system. See [Enabling or disabling system-management settings](#).
 - c. Verify that the user you will log in as has permission to use the FTP interface. See [Managing local users](#).
2. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the destination directory for the log file.
3. Enter:

```
sftp -P port controller-network-address or  
ftp controller-network-address
```

For example:

```
sftp -P 1022 10.235.216.152 or  
ftp 10.1.0.9
```

4. Log in as a user that has permission to use the FTP/SFTP interface.
5. Make sure the client is in binary transfer mode. Enter:
binary
6. Enter:
get logs:heatmap *filename.csv*
where: *filename.csv* is the file that contains the data.

For example:

```
get logs:heatmap IO_density.csv
```

In FTP, wait for the message `Operation Complete` to appear. No messages are displayed in SFTP; instead, the `get` command will return once the download is finished.

7. Quit the FTP/SFTP session.

Using SLP

ME5 Series storage systems support Service Location Protocol (SLP, `srvloc`), which is a service discovery protocol that allows computers and other devices to find services in a LAN without prior configuration. SLP is open for use on all operating systems, and does not require formal licensing.

SLP is based on User Datagram Protocol (UDP) and can use Transmission Control Protocol (TCP) if needed. SLP listens on port 427. When a client, or User Agent (UA), connects to a network, the client queries for Directory Agents (DA) on the network. If no DA responds, the client assumes a DA-less network and sends a multicast UDP query. All Service Agents (SA) that contain query matches will send a UDP answer to the client. If the answer message is too large, the client can repeat the query using TCP.

In a network with DAs, each SA must register all services with a DA. Then the clients will query the DAs, who will respond to the query with its cached SA information.

Through use of DAs, SLP can also scale beyond the local area network to large enterprise, which is an enterprise IT issue. Consult the IETF RFC2165.

When SLP is enabled, the storage system will advertise the interfaces and populate the configuration attributes as shown in the following tables.

You can enable or disable the SLP service in the PowerVault Manager, as described in [Enabling or disabling system-management services](#), or by using the CLI `set protocols` command as described in the CLI Reference Guide.

If the SLP service is enabled, you can test it by using an open source tool, such as `slptool` from www.openslp.org.

Table 19. Interfaces advertised by SLP

Interface (protocol) description	Advertisement string
HTTP	<code>service:api:http</code>
HTTPS	<code>service:api:https</code>
Telnet	<code>service:ui:telnet</code>
SSH	<code>service:ui:ssh</code>
FTP/SFTP (firmware upgrade)	<code>service:firmware-update:ftp/sftp</code>
SNMP	<code>service:api:snmp</code>

Table 20. SLP attributes shown for a storage system

SLP attribute	Corresponding property shown by the CLI <code>show systemdetail</code> command in XML API mode
<code>x-system-name</code>	<code>system-name</code>
<code>x-system-contact</code>	<code>system-contact</code>
<code>x-system-location</code>	<code>system-location</code>
<code>x-system-information</code>	<code>system-information</code>
<code>x-midplane-serial-number</code>	<code>midplane-serial-number</code>
<code>x-vendor-name</code>	<code>vendor-name</code>
<code>x-product-id</code>	<code>product-id</code>
<code>x-product-brand</code>	<code>product-brand</code>
<code>x-wwnn</code>	<code>current-node-wwn</code>
<code>x-platform-type</code>	<code>platform-type</code>
<code>x-bundle-version</code>	no corresponding property
<code>x-build-date</code>	no corresponding property
<code>x-mac-address</code>	no corresponding property
<code>x-top-level-assembly-part-number</code>	no corresponding property

Table 20. SLP attributes shown for a storage system (continued)

SLP attribute	Corresponding property shown by the CLI <code>show systemdetail</code> command in XML API mode
x-top-level-assembly-serial-number	no corresponding property

Administering a log-collection system

A log-collection system receives log data that is incrementally transferred from a storage system for which the managed logs feature is enabled, and is used to integrate the data for display and analysis. For information about the managed logs feature, see [Managed logs](#). Over time, a log-collection system can receive many log files from one or more storage systems. The administrator organizes and stores these log files on the log-collection system. Then, if a storage system experiences a problem that needs analysis, that system's current log data can be collected and combined with the stored historical log data to provide a long-term view of the system's operation for analysis. The managed logs feature monitors the following controller-specific log files:

- Expander Controller (EC) log, which includes EC debug data, EC revisions, and PHY statistics
- Storage Controller (SC) debug log and controller event log
- SC crash logs, which include the SC boot log
- Management Controller (MC) log

Each log-file type also contains system-configuration information.

Topics:

- [How log files are transferred and identified](#)
- [Log file details](#)
- [Storing log files](#)

How log files are transferred and identified

Log files can be transferred to the log-collection system in two ways, depending on whether the managed logs feature is configured to operate in push mode or pull mode.

- Push mode—when log data has accumulated to a significant size, the storage system sends notification events with attached log files through email to the log-collection system. The notification specifies the storage-system name, location, contact, and IP address, and contains a single log segment in a compressed .zip file. The log segment is uniquely named to indicate the log-file type, the date and time of creation, and the storage system. This information is also in the email subject line. The file name format is *logtype_yyyy_mm_dd__hh_mm_ss.zip*.
- Pull mode—when log data has accumulated to a significant size, the system sends notification events using email or SNMP traps to the log-collection system. The notification specifies the information that is transferred such as the storage-system name, location, contact, IP address, and the log-file type (region). The FTP or SFTP interface on the storage system can be used to transfer the appropriate logs to the log-collection system, as described in [Transferring log data to a log-collection system](#).

Log file details


Log files contain the following information.

- SC debug-log records contain date/time stamps of the form *mm/dd hh:mm:ss*.
- SC crash logs (diagnostic dumps) are produced if the firmware fails. Upon restart, such logs are available, and the restart boot log is also included. The four most recent crash logs are retained in the storage system.
- When EC debug logs are obtained, EC revision data and SAS PHY statistics are also provided.
- MC debug logs transferred by the managed logs feature are for five internal components: *appsv*, *mccli*, *logc*, *web*, and *snmpd*. The contained files are log-file segments for these internal components and are numbered sequentially.
- The comments field—used when collecting logs—is limited to 256 characters.

Storing log files

It is recommended to store log files hierarchically by storage-system name, log-file type, and date/time. Then, if historical analysis is required, the appropriate log-file segments can easily be located and can be concatenated into a complete record.

For example, assume that the administrator of a log-collection system has created the following hierarchy for logs from two storage systems named Storage1 and Storage2:



In push mode, when the administrator receives an email with an attached ecdebug file from Storage1, the administrator would open the attachment and unzip it into the ecdebug subdirectory of the Storage1 directory.

In pull mode, when the administrator receives notification that an SC debug log needs to be transferred from Storage2, the administrator would use the FTP or SFTP interface on the storage system to get the log and save it into the scdebug subdirectory of the Storage2 directory.

Settings changed by restoring defaults

This table summarizes the system settings that result from using the `restore defaults` CLI command.

Setting	Value
System information settings	
System name	Uninitialized Name
System contact	Uninitialized Contact
System location	Uninitialized Location
Management protocols settings	
CLI/Telnet	Disabled
CLI/SSH	Enabled
SLP	Enabled
FTP	Disabled
SFTP	Enabled
SNMP	Disabled
WBI/HTTP	Disabled
WBI/HTTPS	Enabled
Debug	Disabled
Ciphers setting	Default cipher strings
CLI session timeout	Preserved
Users	All configured users are deleted and replaced with default user definitions and default settings: User: <code>setup</code> ; Password: press Enter
Tasks and schedules	Preserved
Debug log settings	
Management Controller debug logs	Preserved
Management Controller event logs	Preserved
Storage Controller debug logs	Preserved
Storage Controller event logs	Preserved
Network settings	
Time/date and NTP settings	Preserved
Network IP settings	Preserved
IPv6 network settings	Preserved
DNS management hostname	Preserved
DNS name servers	Preserved
DNS search domains	Preserved

Setting	Value
SNMP settings	
SNMP trap notification level	None
SNMP trap host IPs	0.0.0.0
SNMP read community	public
SNMP write community	private
SMTP settings	
Email notification	Disabled
Email notify filter	None
Email addresses	None
Email server	None
Email domain	None
Email sender	None
Log destination	None
Include logs	Disabled
Alert notification	All
Proxy setting	Cleared
Syslog	
Sylog parameters	Cleared
Syslog settings	Disabled (host IP defaulted to 0.0.0.0)
Alert settings	
Alert condition history	Preserved
Alerts	Preserved
SSL/SSH certificates	Preserved
Host port settings	
FC link speed	Auto
FC topology	Point-to-point
Disk settings	
Disk group metadata	Preserved
Disk spin down	Disabled
Advanced settings	
Disk group background scrub	Enabled
Disk group background scrub interval	24 hours
Partner firmware upgrade	Enabled
Utility priority	High
SMART	Enabled
Dynamic spare configuration	Enabled
Enclosure polling rate	5 seconds
Host control of caching	Disabled

Setting	Value
Sync cache mode	Immediate
Missing LUN response	Not Ready
Controller failure	Disabled
Supercap failure	Enabled
Power supply failure	Disabled
Fan failure	Disabled
Temperature exceeded	Disabled
Partner notify	Disabled
Auto write back	Enabled
Inactive drive spin down	Disabled
Inactive drive spin down delay	15 minutes
Disk background scrub	Disabled
Managed logs	Disabled
Single controller mode	Disabled
Auto stall recovery	Enabled (for failover/failback, not I/O)
Restart on CAPI fail	Enabled
FDE settings	Preserved
Replication settings	
Peer connections	Preserved
Replication sets	Preserved
CHAP records	Preserved
Enclosure settings	
Name	Cleared
Location	Cleared
Rack number	0
Rack position	0.0.0.0
iSCSI port settings	
IP	Preserved
IP version	Preserved
Netmask	Preserved
Gateway	Preserved
Router (IPv6 only)	Preserved
Other iSCSI settings	
CHAP enabled	Preserved
iSNS	Preserved
Jumbo frames	Preserved
Host settings	
Host and initiator nicknames and profiles	Preserved

Setting	Value
Host groups	Preserved
Host port mode	Preserved
Volume settings	
Volume identifying information	Preserved
Volume groups	Preserved
Pool settings	
Thresholds	Preserved
Overcommit	Preserved
Limits and policy	Preserved
Snapshot space thresholds	Preserved
Miscellaneous settings	
CLI parameters	CLI parameters are kept on a per-user basis. All configured users are deleted and replaced with default user definitions and default settings as detailed in the Users section of this table.
Debug log parameters	Each parameter is reset to its default as documented for the <code>set debug-log-parameters</code> CLI command.
Volume snapshot retention priority	Preserved
Volume cache settings	Preserved
Expander PHY settings	Controller module root expander PHY settings are cleared
Volume tier affinity	Preserved
Device identification LED status	Preserved

System configuration limits

The following table lists the system configuration limits for ME5 Series storage systems:

Table 21. ME5 Series system configuration limits

Feature	Value
Enclosures and disks	
Maximum enclosures and disks per system	Supported configurations: <ul style="list-style-type: none"> • 2U12 controller enclosure + nine 2U12 expansion enclosures = 120 • 2U12 controller enclosure + nine 2U24 expansion enclosures = 228 • 2U12 controller enclosure + three 5U84 expansion enclosures = 264 • 2U24 controller enclosure + nine 2U12 expansion enclosures = 132 • 2U24 controller enclosure + nine 2U24 expansion enclosures = 240 • 2U24 controller enclosure + three 5U84 expansion enclosures = 276 • 5U84 controller enclosure + three 5U84 expansion enclosures = 336
Disk groups and pools	
Storage model	Virtual, Linear (mutually exclusive)
Maximum virtual pools per controller module	1
Maximum usable virtual pool size	4016 TiB (4 PB)
Maximum non-ADAPT disk-group	4016 TiB (4 PB)
Maximum disk groups per pool	16
Maximum virtual disk groups per controller module	16
Maximum linear disk groups per controller module	32
Minimum/maximum disks per virtual disk group	<ul style="list-style-type: none"> • NRAID (non-RAID): 1/1 (read cache only) • RAID 0: 2/2 (read cache only) • RAID 1: 2/2 • RAID 5: 3/16 • RAID 6: 4/16 • RAID 10: 4/16 • ADAPT: 12/128
Minimum/maximum disks per linear disk group	<ul style="list-style-type: none"> • RAID 0: 2/16 • RAID 1: 2/2 • RAID 5: 3/16 • RAID 6: 4/16 • RAID 10: 4/16 • ADAPT:12/128
Maximum dedicated spares per linear disk group	4
Maximum global spares per system	64
Maximum ADAPT groups per controller module	4
Maximum ADAPT single disk size	64 TiB

Table 21. ME5 Series system configuration limits (continued)

Feature	Value
Maximum ADAPT disk group size	1.5 PiB
ADAPT stripe width (data+parity)	8+2, 16+2
Volumes, initiators, hosts, and mapping	
Maximum virtual volumes per system	1024
Maximum linear volumes per system	1024 (512 recommended)
Maximum linear volume (LUN) size	Only limited by 64-bit addressing, so 8 ZiB with 512-byte sectors
Maximum virtual volume (LUN) size	128 TiB
Maximum mappable volumes (LUNs) per disk group	128
Maximum mappable virtual volumes (LUNs) per pool	512
Maximum mappable linear volumes (LUNs) per pool	128
Maximum mappable volumes (LUNs) per controller module	512
Maximum virtual volumes per pool	1024 (512 base volumes and 512 snapshots)
Maximum linear volumes per pool	1024
Maximum virtual volumes per volume group	1024
Maximum volume groups per system	256
Maximum volumes per replication volume group	16
Maximum initiators per controller module	4096
Maximum initiators per system	8192
Maximum initiators per volume	128
Maximum initiators per host	128
Maximum hosts per host group	256
Maximum host groups per system	32
Maximum commands per LUN (preferred path)	1024 per port
Maximum queue depth per host port	1024
Maximum FC host-port link speed	32 Gb
Maximum iSCSI host-port link speed	25 Gb
Maximum SAS host-port link speed	12 Gb
Virtual volume snapshots	
Maximum snapshots per system, with a license	Not applicable
Maximum snapshots per pool (net usable)	512
Maximum base volumes per system	1024
Maximum base snapshots per base volume	254 in the volume snapshot tree
Maximum mappable snapshots per system	1024
Virtual volume replication	
Maximum number of peer connections per system	4
Maximum number of replicated volumes per system	32
Maximum number of replication sets for a volume	1

Table 21. ME5 Series system configuration limits (continued)

Feature	Value
Maximum number of volumes for a replicated volume group	16, if no other volumes belong to a replication set
Minimum replication schedule interval	30 minutes
Miscellaneous	
Maximum SCSI reservations per system	1024
Maximum SCSI reservations per LUN	1
Maximum SCSI registrations per system	32768
Maximum SCSI reservations per virtual LUN	4096
Maximum SCSI reservations per linear LUN	<ul style="list-style-type: none">• FC: 85• iSCSI: 85-128 depending on IQN length• SAS: 85

Best practices

This appendix describes best practices for configuring and provisioning a storage system.

Topics:

- [Pool setup](#)
- [RAID selection](#)
- [Disk count per RAID level](#)
- [Disk groups in a pool](#)
- [Tier setup](#)
- [Multipath configuration](#)

Pool setup

In a storage system with two controller modules, try to balance the workload of the controllers. Each controller can own one virtual pool. Having the same number of disk groups and volumes in each pool will help balance the workload, increasing performance.

RAID selection

A pool is created by adding disk groups to it. Disk groups are based on RAID technology.

The following table describes the characteristics and use cases of each RAID level:

Table 22. RAID level characteristics and use cases

RAID level	Protection	Performance	Capacity utilization	Application use cases	Suggested disk speed
RAID 1/RAID 10	Protects against up to one disk failure per mirror set	Great random I/O performance	Poor: 50% fault tolerance capacity loss	Databases, OLTP, Exchange Server	10K, 7K
RAID 5	Protects against up to one disk failure per RAID set	Good sequential I/O performance, good random read I/O performance, moderate random write performance	Great: One-disk fault tolerance capacity loss	Big data, media and entertainment (ingest, broadcast, and post production)	10K, lower capacity 7K
RAID 6	Protects against up to two disk failures per RAID set	Good sequential I/O performance, good random read I/O performance, moderate random write performance (lower than RAID5)	Moderate: Two-disk fault tolerance capacity loss	Archive, parallel distributed file system	High capacity 7K
ADAPT	Distributed erasure coding with dual disk failure protection supports 16+2 or 8+2 stripe width	Very fast rebuilds, no spare disks (built-in spare capacity), large storage pools, simplified initial	0% of capacity used for fault tolerance for 8+2; 11.1% of capacity used for fault tolerance for 16+2	Provides flexible storage, well suited or most workloads other than those using very few disks, or requiring a high number of writes	High capacity 7K

Table 22. RAID level characteristics and use cases (continued)

RAID level	Protection	Performance	Capacity utilization	Application use cases	Suggested disk speed
		deployment and expansion			

Disk count per RAID level

The controllers allocate virtual volume storage in 4-MiB pages, which are referenced paged tables in memory. There is a sequential write performance penalty when RAID-5 or RAID-6 disk groups are used in a virtual pool and the stripe size of the disk group doesn't align well with the 4-MiB page.

- Example 1: Consider a RAID-5 disk group with five disks. The equivalent of four disks provide usable capacity, and the equivalent of one disk is used for parity (parity is distributed among disks). The four disks providing usable capacity are the data disks and the one disk providing parity is the parity disk. In reality, the parity is distributed among all the disks, but conceiving of it in this way helps with the example.
 - ⓘ **NOTE:** The number of data disks is a power of two (2, 4, and 8). The controller will use a 512-KiB stripe unit size when the data disks are a power of two. This results in a 4-MiB page being evenly distributed across two stripes. This is good for performance. Even better would be a 9-disk RAID-5 disk group with 8 total data disks, allowing one page to fit in exactly one stripe.
- Example 2: Consider a RAID-5 disk group with six disks. The equivalent of five disks now provide usable capacity. Assume the controller again uses a stripe unit of 512-KiB. When a 4-MiB page is pushed to the disk group, one stripe will contain a full page, but the controller must read old data and old parity from two of the disks in combination with the new data in order to calculate new parity. This is known as a read-modify-write, and it is a performance killer with sequential workloads. In essence, every page push to a disk group would result in a read-modify-write.

To mitigate this issue, the controllers use a stripe unit of 64-KiB when a RAID-5 or RAID-6 disk group isn't created with a power-of-two data disks. This results in many more full-stripe writes, but at the cost of many more I/O transactions per disk to push the same 4-MiB page.

The following table shows recommended disk counts for RAID-6 and RAID-5 disk groups. Each entry specifies the total number of disks and the equivalent numbers of data and parity disks in the disk group. Note that parity is actually distributed among all the disks.

Table 23. Recommended disk group sizes

RAID level	Total disks	Data disks (equivalent)	Parity disks (equivalent)
RAID 6	4	2	2
	6	4	2
	10	8	2
RAID 5	3	2	1
	5	4	1
	9	8	1

To ensure best performance with sequential workloads and RAID-5 and RAID-6 disk groups, use a power-of-two data disks.

Disk groups in a pool

For better efficiency and performance, use similar disk groups in a pool.

- Disk count balance: For example, with 20 disks, it is better to have two 8+2 RAID-6 disk groups than one 10+2 RAID-6 disk group and one 6+2 RAID-6 disk group.
- RAID balance: It is better to have two RAID-5 disk groups than one RAID-5 disk group and one RAID-6 disk group.
- In terms of the write rate, due to wide striping, tiers and pools are as slow as their slowest disk groups.
- All disks in a tier should be the same type.

Tier setup

In general, it is best to have two tiers instead of three tiers. The highest tier will nearly fill before using the lowest tier. The highest tier must be 95% full before the controller will evict cold pages to a lower tier to make room for incoming writes.

Typically, you should use tiers with SSDs and 10K disks, or tiers with SSDs and 7K disks. An exception may be if you need to use both SSDs and faster spinning disks to hit a combination of price for performance, but you cannot hit your capacity needs without the 7K disks; this should be rare.

Multipath configuration

ME5 Series storage systems comply with the SCSI-3 standard for Asymmetrical Logical Unit Access (ALUA).

ALUA compliant storage systems provide optimal and non-optimal path information to the host during device discovery. To implement ALUA, you must configure your servers to use multipath I/O (MPIO).

Use one of the following procedures to enable MPIO.

Enabling MPIO on Windows

1. Start Server Manager if it is not already running.
2. In the Manage menu, select **Add Roles and Features**.
3. In the Add Roles and Features Wizard, select **Role-based or Feature Based Installation**.
4. Click **Next**.
5. Select the server from the pool and then click **Next**.
6. Click **Next** again to go to the feature selection window.
7. Select the **Multipath IO** check box and then click **Next**.
8. Click **Install**.
9. When the installation is complete, click **Close**.
10. In the Server Manager Tools menu, select **MPIO**.
11. Select the **Discover Multi-Paths** tab.
12. Select check boxes for the devices you want to support, and then click **Add**.
13. When prompted, reboot the system.

When the reboot is complete, MPIO is ready to use.

Enabling MPIO on Linux

This procedure is a general procedure, and does not address specific nuances of different Linux operating systems.

1. Run the following command to ensure that the multipath daemon is installed and set to start at run-time:

```
systemctl status multipathd
```

2. Ensure the correct entries exist in the `/etc/multipath.conf` file on each OSS/MDS host. Create a separate device entry for the ME5 Series storage system. The following table specifies four attributes that should be set. Run the following command to obtain the exact vendor and product ID values:

```
multipath -v3
```

Attribute	Value
prio	alua
failback	immediate
vendor	<i>vendor-name</i>
product	<i>product-ID</i>

Retain the default settings for the remaining attributes. For example:

```
devices
device {
    vendor "DellEMC"
    product "ME5"
    path_checker "tur"
    hardware_handler "1 alua"
    prio "alua"
    failback immediate
    rr_weight "uniform"
    path_selector "service-time 0"
}
```

3. Run the following command to reload the multipath.conf file:

```
systemctl restart multipathd
```

4. Run the following command to determine if the multipath daemon used ALUA to obtain the optimal/non-optimal paths:

```
multipath -v3 | grep alua
```

You should see output stating that ALUA was used to configure the path priorities. For example:

```
Oct 01 14:28:43 | sdb: prio = alua (controller setting) Oct 01 14:28:43 | sdb: alua prio = 130
```

Glossary of terms

The following table lists definitions of the terms used in ME5 Series publications:

Table 24. Glossary of ME5 Series terms

Term	Definition
2U12	An enclosure that is two rack units in height and can contain 12 disks.
2U24	An enclosure that is two rack units in height and can contain 24 disks.
5U84	An enclosure that is five rack units in height and can contain 84 disks.
ADAPT	A RAID-based data protection level that maximizes flexibility, provides built in spare capacity, and allows for very fast rebuilds, large storage pools, and simplified expansion.
AES	Advanced Encryption Standard.
AFA	All-flash array. A storage system that uses only SSDs, without tiering.
all-flash array	All-flash array (AFA). A storage system that uses only SSDs, without tiering.
allocated page	A page of virtual pool space that has been allocated to a volume to store data.
allocation rate	The rate, in pages per minute, at which a virtual pool is allocating pages to its volumes because they need more space to store data.
array	See storage system.
ASC/ASCQ	Additional Sense Code/Additional Sense Code Qualifier. Information on sense data returned by a SCSI device.
atomic write	A write-optimization mode for volume cache that guarantees if a failure (such as I/O being aborted or a controller failure) interrupts a data transfer between a host system and the storage system, controller cache will contain either all the old data or all the new data, not a mix of old and new data. This option has a slight performance cost because it maintains a secondary copy of data in cache so that if a data transfer is not completed, the old cache data can be restored.
auto-write-through	Auto-write-through (AWT). A setting that specifies when the RAID controller cache mode automatically changes from write-back to write-through.
automated tiered storage	A virtual-storage feature that automatically uses the appropriate tier of disks to store data based on how frequently the data is accessed. This enables higher-cost, higher-speed disks to be used only for frequently needed data, while infrequently needed data can reside in lower-cost, lower-speed disks.
available disk	A disk that is not a member of a disk group, is not configured as a spare, and is not in the leftover state. It is available to be configured as a part of a disk group or as a spare.
AWT	Auto-write-through. A setting that specifies when the RAID controller cache mode automatically changes from write-back to write-through.
base volume	A virtual volume that is not a snapshot of any other volume, and is the root of a snapshot tree.
canister	Synonym for IOM.
CAPI	Configuration Application Programming Interface. A proprietary protocol used for communication between the Storage Controller and the Management Controller in a controller module. CAPI is always enabled.
CHAP	Challenge-Handshake Authentication Protocol.
chassis	The sheet metal housing of an enclosure.

Table 24. Glossary of ME5 Series terms (continued)

Term	Definition
child volume	The snapshot of a parent volume in a snapshot tree. See parent volume.
chunk size	The amount of contiguous data that is written to a disk group member before moving to the next member of the disk group.
compatible disk	A disk that can be used to replace a failed member disk of a disk group because it has at least the same capacity as, and is of the same type (enterprise SAS, for example) as, the disk that failed.
controller A (or B)	A short way of referring to controller module A (or B).
controller enclosure	An enclosure that contains one or two controller modules.
controller module	A FRU that contains the following subsystems and devices: a Storage Controller processor; a Management Controller processor; a SAS expander and Expander Controller processor; management interfaces; cache protected by a super capacitorpack and non-volatile memory; host, expansion, network, and service ports; and midplane connectivity.
CPLD	Complex programmable logic device.
CRC	Cyclic Redundancy Check.
CRU	Customer-replaceable unit. See customer FRU.
CSV	Comma-separated values. A format to store tabular data in plain-text form.
customer FRU	A product module that can be ordered as a SKU and replaced in an enclosure by customers or by qualified service personnel, without having to send the enclosure to a repair facility.
DAS	Direct Attached Storage. A dedicated storage device that connects directly to a host without the use of a switch.
deallocation rate	The rate, in pages per minute, at which a virtual pool is deallocating pages from its volumes because they no longer need the space to store data.
dedicated spare	A disk that is reserved for use by a specific linear disk group to replace a failed disk. See also available disk, compatible disk, dynamic spare, global spare.
DES	Data Encryption Standard.
DHCP	Dynamic Host Configuration Protocol. A network configuration protocol for hosts on IP networks.
disk group	A group of disks that is configured to use a specific RAID level and provides storage capacity for a pool. The number of disks that a disk group can contain is determined by its RAID level.
disk spin down	Disk spin down (DSD). A power-saving feature for spinning disks that monitors disk activity and spins down inactive disks based on user-selectable policies.
DNS	Domain Name System.
drain	The automatic movement of active volume data from a virtual disk group to other disk-group members within the same pool.
drawer	In a 5U84 enclosure, one of two FRUs that each holds 42 disks
drive enclosure	See expansion enclosure. See also EBOD, JBOD.
drive spin down	Disk spin down (DSD). A power-saving feature for spinning disks that monitors disk activity and spins down inactive disks based on user-selectable policies
DRM	Disaster recovery management. Storage-system firmware features that, when the Site Replication Adapter (SRA) feature is enabled, support the use of VMware Site Recovery Manager to automate disaster-recovery failover and failback tasks. See also SRA.
DSD	Disk spin down (DSD). A power-saving feature for spinning disks that monitors disk activity and spins down inactive disks based on user-selectable policies.
DSP	Digital signal processor.

Table 24. Glossary of ME5 Series terms (continued)

Term	Definition
dual-port disk	A disk that is connected to both controllers so it has two data paths, achieving fault tolerance.
dynamic spare	An available compatible disk that is automatically assigned, if the dynamic spares option is enabled, to replace a failed disk in a disk group with a fault-tolerant RAID level.
EBOD	Expanded Bunch of Disks. Expansion enclosure attached to a controller enclosure.
EC	Expander Controller. A processor (located in the SAS expander in each controller module and expansion module) that controls the SAS expander and provides SES functionality. See also EMP.
EEPROM	Electrically erasable programmable ROM.
eMMC	Embedded MultiMedia card. Also referred to as memory card, non-volatile memory.
EMP	Enclosure management processor. An Expander Controller subsystem that provides SES data such as temperature, power supply and fan status, and the presence or absence of disks.
enclosure	A physical storage device that contains I/O modules, disk drives, and other FRUs.
enclosure management processor	Enclosure management processor (EMP). An Expander Controller subsystem that provides SES data such as temperature, power supply and fan status, and the presence or absence of disks
ESD	Electrostatic discharge.
Expander Controller	See EC.
expansion enclosure	An enclosure that contains one or two expansion modules. Expansion enclosures can be connected to a controller enclosure to provide additional storage capacity. See also EBOD, JBOD.
expansion module	A FRU that contains the following subsystems and devices: a SAS expander and Expander Controller processor; host, expansion, and service ports; and midplane connectivity. In an expansion enclosure, the upper expansion module is designated A and the lower one is designated B.
failback	Synonym for recovery.
failover	In an active-active configuration, failover is the act of temporarily transferring ownership of controller resources from an offline controller to its partner controller, which remains operational. The resources include pools, volumes, cache data, host ID information, and LUNs and WWNs. See also recovery.
fan module	The fan FRU used in 5U84 enclosures. There are five in each enclosure, separate from the PSUs.
FC	Fibre Channel interface protocol.
FC-AL	Fibre Channel Arbitrated Loop. The FC topology in which devices are connected in a one-way loop.
FDE	Full Disk Encryption. A feature that secures all the user data on a storage system. See also lock key, passphrase, repurpose, SED.
FPGA	Field-programmable gate array. An integrated circuit designed to be configured after manufacturing.
FQDN	Fully qualified domain name.
FRU	Field-replaceable unit. See service FRU.
Full Disk Encryption	Full Disk Encryption (FDE). A feature that secures all the user data on a storage system. See also lock key, passphrase, repurpose, SED.
GEM	Generic Enclosure Management. The firmware responsible for managing enclosure electronics and environmental parameters. GEM is used by the Expander Controller.

Table 24. Glossary of ME5 Series terms (continued)

Term	Definition
global spare	A compatible disk that is reserved for use by any disk group with a fault-tolerant RAID level to replace a failed disk.
HBA	Host bus adapter. A device that facilitates I/O processing and physical connectivity between a host and the storage system.
HDD	Hard disk drive
host	A user-defined object that represents a server to which the storage system is attached, and is used to define a mapping relationship to storage.
host group	A user-defined group of hosts for ease of management, such as for mapping operations.
host port	A port on a controller module that interfaces to a host computer, either directly or through a network switch.
I/O Manager	An SNMP MIB term for a controller module.
I/O module	Input/output module (I/O module, IOM). An IOM can be either a controller module or an expansion module
initiator	An external port to which the storage system is connected. The external port may be a port in an I/O adapter in a server, or a port in a network switch.
IOM	Input/output module, or I/O module. An IOM can be either a controller module or an expansion module.
IOPS	I/O operations per second.
IQN	iSCSI Qualified Name.
iSCSI	Internet SCSI interface protocol
ISE	Instant Secure Erase. A feature designed to protect data on hard disk drives by instantly resetting the drive back to factory settings and changing the encryption key so that any data remaining on the drive is cryptographically erased. This means all data on the drive is permanently and instantly unreadable, as needed
iSNS	Internet Storage Name Service.
JBOD	"Just a bunch of disks." An expansion enclosure attached to a server.
LBA	Logical block address. The address used for specifying the location of a block of data.
leftover	The state of a disk that the system has excluded from a disk group because the timestamp in the metadata on the disk is older than the timestamp of other disks in the disk group, or because the disk was not detected during a rescan. A leftover disk cannot be used in another disk group until the metadata on the disk is cleared. For information and cautions about doing so, see documentation topics about clearing disk metadata.
LFF	Large form factor.
linear	The storage-class designation for logical components such as volumes that store user data in sequential, fully allocated physical blocks, using a fixed (static) mapping between the logical data presented to hosts and the physical storage where it is stored.
LIP	Loop Initialization Primitive. An FC primitive used to determine the loop ID for a controller.
lock key	A system-generated value that manages the encryption and decryption of data on FDE-capable disks. See also FDE, passphrase.
loop	Fibre Channel Arbitrated Loop. The FC topology in which devices are connected in a one-way loop.
LUN	Logical Unit Number. A number that identifies a mapped volume to a host system.
MAC address	Media Access Control Address. A unique identifier assigned to network interfaces for communication on a network.

Table 24. Glossary of ME5 Series terms (continued)

Term	Definition
Management Controller	Management Controller (MC). A processor (located in a controller module) that is responsible for human-computer interfaces, such as a WBI, and computer-computer interfaces, such as SNMP, and interacts with the Storage Controller.
map	Settings that specify whether a volume is presented as a storage device to a host, and how the host can access the volume. Mapping settings include an access type and a LUN that identifies the volume to the host.
MC	Management Controller. A processor (located in a controller module) that is responsible for human-computer interfaces, such as the PowerVault Manager, and computer-computer interfaces, such as SNMP, and interacts with the Storage Controller.
metadata	Data in the first sectors of a disk that stores disk-, disk-group-, and volume-specific information including disk group membership or spare identification, disk group ownership, volumes and snapshots in the disk group, host mapping of volumes, and results of the last media scrub.
MIB	Management Information Base. A database used for managing the entities in SNMP. midplane The printed circuit board to which components connect in the middle of an enclosure. mount To enable access to a volume from a host OS. See also host, map/mapping, volume.
midplane	The printed circuit board to which components connect in the middle of an enclosure.
mount	To enable access to a volume from a host OS. See also host, map/mapping, volume.
network port	The Ethernet port on a controller module through which its Management Controller is connected to the network.
NRAID	Non-RAID, non-striped mapping to a single disk.
NTP	Network time protocol.
OID	Object Identifier. In SNMP, an identifier for an object in a MIB.
orphan data	See unwritable cache data.
overcommit	A setting that controls whether a virtual pool is allowed to have volumes whose total size exceeds the physical capacity of the pool.
overcommitted	The amount of storage capacity that is allocated to virtual volumes exceeds the physical capacity of the storage system.
page	A range of contiguous LBAs in a virtual disk group.
paged storage	A method of mapping logical host requests to physical storage that maps the requests to virtualized “pages” of storage that are in turn mapped to physical storage. This provides more flexibility for expanding capacity and automatically moving data than the traditional, linear method in which requests are directly mapped to storage devices. Paged storage is also called virtual storage.
parent volume	A virtual volume that has snapshots (can be either a base volume or a base snapshot volume). The parent of a snapshot is its immediate ancestor in the snapshot tree.
partner firmware update	Partner firmware update (PFU). The automatic update of the partner controller when the user updates firmware on one controller.
passphrase	A user-created password that allows users to manage lock keys in an FDE-capable system. See also FDE, lock key.
PCB	Printed circuit board.
PCBA	Printed circuit board assembly.
PCM	Power and cooling module FRU. A power supply module that includes an integrated fan. See also PSU.
PDU	Power distribution unit. The rack power-distribution source to which a PCM or PSU connects.

Table 24. Glossary of ME5 Series terms (continued)

Term	Definition
peer connection	The configurable entity defining a peer-to-peer relationship between two systems for the purpose of establishing an asynchronous replication relationship. See also peer system.
peer system	A remote storage system that can be accessed by the local system and is a candidate for asynchronous replications. Both systems in a peer connection are considered peer systems to each other, and they both maintain a peer connection with the other. Asynchronous replication of volumes may occur in either direction between peer systems configured in a peer connection.
PFU	Partner firmware update. The automatic update of the partner controller when the user updates firmware on one controller.
PGR	Persistent group reservations.
PHY	One of two hardware components that form a physical link between devices in a SAS network that enables transmission of data.
point-to-point	Fibre Channel Point-to-Point topology in which two ports are directly connected.
pool	container for volumes that is composed of one or more virtual disk groups.
POST	Power-on self test. Tests that run immediately after a device is powered on.
PowerVault Manager	PowerVault Manager is the web-browser interface (WBI), the web application that is embedded in each controller module and is the primary management interface for the storage system
primary system	The storage system that contains the primary volume of a replication set.
primary volume	The volume that is the source of data in a replication set and that can be mapped to hosts. The primary volume exists in a primary (linear storage) or pool (virtual storage) in the primary storage system.
PSU	Power supply unit FRU.
RAID head	Synonym for controller enclosure.
RBOD	“RAID bunch of disks.” See controller enclosure.
read cache	A special virtual disk group, comprised of SSDs, that can be added to a pool for the purpose of speeding up read access to data stored on spinning disks elsewhere in the pool.
recovery	In an active-active configuration, recovery is the act of returning ownership of controller resources to a controller (which was offline) from its partner controller. The resources include volumes, cache data, host ID information, and LUNs and WWNs. See also failover.
replication	Asynchronous replication of block-level data from a volume in a primary system to a volume in a secondary system by creating an internal snapshot of the primary volume and copying the snapshot data to the secondary system via Fibre Channel or iSCSI links.
replication set	A container that houses the infrastructure upon which replications are performed. It defines a relationship between a primary and secondary volume for the purposes of maintaining a remote copy of the primary volume on a peer system. See primary volume, secondary volume.
replication set failover	The secondary system of a replication set has allowed direct access to the secondary volume or volume group because the primary system is not operational. In this state no replications will occur, even if the primary system becomes operational and communication is restored. The secondary volume can be mapped and accessed for use, including rollback to the contents of any manually created or snapshot-history snapshot.
replication snapshot history	As part of handling a replication, the replication set will automatically take a snapshot of the primary and/or secondary volume, thereby creating a history of data that has been replicated over time. This feature can be enabled for a secondary volume or for a primary volume and its secondary volume, but not for a volume group.
repurpose	A method by which all data on a system or disk is erased in an FDE-capable system. Repurposing unsecures the system and disks without needing the correct passphrase. See also FDE, passphrase.

Table 24. Glossary of ME5 Series terms (continued)

Term	Definition
SAS	Serial Attached SCSI.
SATA	Serial ATA.
SC	Storage Controller. A processor (located in a controller module) that is responsible for RAID controller functions. The SC is also referred to as the RAID controller. See also EC, MC.
secondary system	The storage system that contains the secondary volume of a replication set. See also replication set, primary system.
secondary volume	The volume that is the destination for data in a replication set and that is not accessible to hosts. The secondary volume exists in a secondary (linear storage) or pool (virtual storage) in a secondary storage system.
secret	For use with CHAP, a password that is shared between an initiator and a target to enable authentication.
SED	Self-encrypting drive. A disk drive that provides hardware-based data encryption and supports the Full Disk Encryption feature. See also FDE.
SEEPROM	Serial electrically erasable programmable ROM. A type of nonvolatile (persistent if power removed) computer memory used as FRU ID devices.
service FRU	A product module that can be replaced in an enclosure by qualified service personnel only, without having to send the enclosure to a repair facility.
service tag	An alphanumeric string that uniquely identifies the product and is used to appropriately route support cases.
SES	SCSI Enclosure Services. The protocol that allows the initiator to communicate with the enclosure using SCSI commands.
SFF	Small form factor.
SFTP	SSH File Transfer Protocol. A secure secondary interface for installing firmware updates, downloading logs, and installing security certificates and keys. All data sent between the client and server will be encrypted.
SHA	Secure Hash Algorithm.
shelf	Synonym for enclosure.
sideplane	A printed circuit board to which components connect longitudinally within an enclosure.
SLP	Service Location Protocol. Enables computers and other devices to find services in a local area network without prior configuration.
SMART	Self-Monitoring Analysis and Reporting Technology. A monitoring system for disk drives that monitors reliability indicators for the purpose of anticipating disk failures and reporting those potential failures.
snapshot	A point-in-time copy of the data in a source volume that preserves the state of the data as it existed when the snapshot was created. Data associated with a snapshot is recorded in the source volume. A snapshot can be mapped and written to. Snapshots that can be mapped to hosts are counted against the snapshot-license limit, whereas transient and unmappable snapshots are not.
snapshot tree	A group of virtual volumes that are interrelated due to creation of snapshots. Since snapshots can be taken of existing snapshots, volume inter-relationships can be thought of as a "tree" of volumes. A tree can be 254 levels deep. See also base volume, child volume, parent volume, source volume.
source volume	A volume that has snapshots. Used as a synonym for parent volume.
SRA	Storage Replication Adapter. A host-based software component that allows VMware Site Recovery Manager to manage the disaster recovery management (DRM) features of the storage-system firmware, automating disaster-recovery failover and failback tasks. The SRA uses the CLI XML API to control the storage system. See also DRM.

Table 24. Glossary of ME5 Series terms (continued)

Term	Definition
SSD	Solid-state drive.
SSH	Secure Shell. A network protocol for secure data communication.
SSL	Secure Sockets Layer. A cryptographic protocol that provides security over the internet.
standard volume	A volume that can be mapped to initiators and presented as a storage device to a host system, but is not enabled for snapshots.
Storage Controller	Storage Controller. A processor (located in a controller module) that is responsible for RAID controller functions. The SC is also referred to as the RAID controller. See also EC, MC.
storage system	A controller enclosure with at least one connected expansion enclosure. Product documentation and interfaces use the terms storage system and system interchangeably.
syslog	A protocol for sending event messages across an IP network to a logging server. This feature supports User Datagram Protocol (UDP) but not Transmission Control Protocol (TCP).
TCP	Transmission Control Protocol.
thin provisioning	A virtual-storage feature that allows actual storage for a virtual volume to be assigned as data is written, rather than storage being assigned immediately for the eventual size of the volume. This allows the storage administrator to overcommit physical storage, which in turn allows the connected host system to operate as though it has more physical storage available than is actually allocated to it. When physical resources fill up, the storage administrator can add storage capacity on demand.
tier	A homogeneous group of disks, typically of the same capacity and performance level, that comprise one or more virtual disk groups in the same pool. Tiers differ in their performance, capacity, and cost characteristics, which forms the basis for the choices that are made with respect to which data is placed in which tier. The predefined tiers are: Performance, which uses SSDs; Standard, which uses enterprise-class spinning SAS disks; Archive, which uses midline spinning SAS disks.
tier migration	The automatic movement of blocks of data, associated with a single virtual volume, between tiers based on the access patterns that are detected for the data on that volume.
tray	Synonym for enclosure.
UDP	User Datagram Protocol.
ULP	Unified LUN Presentation. A RAID controller feature that enables a host system to access mapped volumes through any controller host port. ULP incorporates ALUA extensions.
undercommitted	The amount of storage capacity that is allocated to volumes is less than the physical capacity of the storage system.
unmount	To remove access to a volume from a host OS. Synonyms include unpresent and unmap.
unwritable cache data	Cache data that has not been written to disk and is associated with a volume that no longer exists or whose disks are not online. If the data is needed, the volume disks must be brought online. If the data is not needed it can be cleared, in which case it will be lost and data will differ between the host system and disk. Unwritable cache data is also called orphan data.
UPS	Uninterruptible power supply.
UTC	Coordinated Universal Time.
UTF-8	UCS transformation format - 8-bit. A variable-width encoding that can represent every character in the Unicode character set used for the PowerVault Manager and CLI interfaces.
virtual	The storage-class designation for logical components such as volumes that use paged-storage technology to virtualize data storage. See paged storage.
volume	A logical representation of a fixed-size, contiguous span of storage that is presented to host systems for the purpose of storing data.
volume copy	An independent copy of the data in a linear volume. The capability to copy volumes makes use of snapshot functionality.

Table 24. Glossary of ME5 Series terms (continued)

Term	Definition
volume group	A user-defined group of volumes for ease of management, such as for host-attachment operations.
VPD	Vital Product Data. Data held on an EEPROM in an enclosure or FRU that is used by GEM to identify and control the component.
WBI	PowerVault Manager is the web-browser interface (WBI), the web application that is embedded in each controller module and is the primary management interface for the storage system.
WWN	World Wide Name. A globally unique 64-bit number that identifies a device used in storage technology.
WWNN	World Wide Node Name. A globally unique 64-bit number that identifies a device.
WWPN	World Wide Port Name. A globally unique 64-bit number that identifies a port.