Exos® CORVAULT™
Storage Management Guide

Abstract

This guide provides information about managing a Seagate Exos CORVAULT storage system by using its web interface, the Storage Management Console (SMC).

Part Number: 205042000-01
Revision: B
Published: June 2024

Contents

Tables

1 Getting started

The Storage Management Console (SMC) is a web-based application for configuring, monitoring, and managing the storage system. The SMC 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 SMC. 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 SMC, 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*.

TIP The Quick Configuration Guide (QCG) is a pamphlet included in the enclosure shipkit. It provides a quick reference for using management interfaces to configure a new storage system, and complete preboarding and onboarding.

Product features

The SMC gives you access to many features that help you manage the storage system. Some of the key features include:

- **Storage system**: two different enclosure configurations are supported. The SAS HDD configuration ships with 106 drives and the SATA configuration ships with 96 drives. Each storage configuration provides intuitive system setup. See the Hardware Installation and Maintenance Guide for more detail.
- **Storage system setup**: the capability to initially connect to a system using the SMC, which employs intuitive preboarding and onboarding steps to guide you through initial setup of your storage, as described in "Configuring and provisioning a new storage system" on the next page.
- **ADAPT data protection**: RAID-based data protection level that emphasizes efficiency, as described in "ADAPT" on page 20. Interleaving can provide improved performance as described in "ADAPT interleaved volumes" on page 21.
- **Maximum disk size**: the maximum single disk size supported for drives used in storage enclosures is 26TB. See also "System configuration limits" on page 94.
- **SPEAR** (Stripe Parity Error Assessment on Read): the ability to check the parity of blocks in stripes with each read access, before returning data to the host, as described in "SPEAR" on page 21.
- **Autonomous Drive Regeneration** (**ADR**): the ability to repair a spinning disk that has become unusable due to a single head failure, and then seamlessly add it back to the ADAPT disk group when the process is complete, as described in "Autonomous drive regeneration" on page 24.
- **Update firmware**: the capability to notify users of available firmware updates to controller modules and disk modules with newer/compatible firmware versions as they become available, as described in "Updating firmware" on page 32.
- **Alerts**: a robust storage enclosure health and notification system designed to identify actionable conditions and promote best practices, as described in "Alerts panel" on page 42.
- **LDAP integration**: the capability to use the external Lightweight Directory Access Protocol services on Windows systems for user authentication and authorization, as described in "LDAP" on page 34.
- **SSDs**: the ability to use solid-state drives to enhance storage system performance, as described in "SSDs" on page 25.
- **Linear storage**: a storage model that maps logical components to physical media, as described in "Linear storage" on page 16.
- **Hedged reads**: reduces latency for read requests by using RAID parity information to reconstruct requested data, as described in "Hedged reads" on page 29.

- **IPv6 support**: the capability for the storage system to support IPv6 (Internet Protocol version 6) functionality (in addition to IPv4), as described in "Configuring controller network ports" on page 51.
- **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 related to the storage provisioning model, including disks, storage pools, volumes, and enclosures.

  The public API 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 SNIA Swordfish, see https://www.snia.org/forums/smi/swordfish.
  - For technical information about Seagate's implementation of the Redfish and Swordfish application programming interfaces—and how to use them—see the *Seagate Redfish/Swordfish API Guide*.
  - The base URL for accessing the Redfish API capability is "`https://<controller-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.

Configuring and provisioning a new storage system

When you connect to the system for the first time, a wizard in the SMC guides you through the first-time setup of your system. This process is referred to as preboarding and onboarding. During preboarding you are led through steps to prepare the system for use and are prompted to do the following:

- Create a username and password (once complete, you will be logged into the system as this user)
- Update firmware

NOTE The user created during the preboarding process will have managing capabilities and will be able to change system settings.

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

- Configuring system settings:
  - Network settings (IPv4, IPv6, DNS, including setting a proxy)
  - Date and time (NTP)
  - User definitions (local, LDAP, SNMPv3)
  - Notifications (email, SNMP, syslog)
- Configuring storage settings (select one option):

  TIP To simplify storage system configuration—and make use of key features—choose one of the preconfigured ADAPT protection levels from the wizard, based on your storage requirements. The first four options in the following list pertain to enclosures equipped with SAS drives. The SATA configuration options pertain to enclosures equipped with variable-capacity SATA drives.

  - **Highest Capacity**: ADAPT + interleaved, 16+2 stripe width (default)
  - **Highest Sequential Performance**: ADAPT + non-interleaved, 8+2 stripe width

  - **Highest Reliability**: ADAPT + non-interleaved, 16+3 stripe width
  - **Manual** (advanced)
  - SATA Configuration 1 (ADAPT + interleaved, 16+3 stripe width)
  - SATA Configuration 2 (ADAPT + non-interleaved, 16+3 stripe width)

  Each storage configuration option includes a general description. Additional technical details are provided for the **Highest Capacity**, **Highest Reliability**, and **Highest Sequential Performance** options. Related to these settings, see "Scripting reference for CLI setup commands" on page 96.

  If working with variable-capacity SATA drives—pertaining to SATA configurations—a general description and additional technical details are provided for each of the SATA configuration options within the SMC.

  NOTE Prerequisites for **Highest Capacity**, **Highest Reliability**, and **Highest Sequential Performance** options:

  - Storage system is fully populated with identical working disks
  - Storage system is not configured (no existing volumes or disk groups)

  If the above requirements are not met, the storage system configuration is set to **Manual**, and an informational message is provided. See also "Linear disk group configuration" on page 17.

- Provisioning storage:
  - Creating hosts and host groups (naming initiators, assigning initiators to hosts, creating a single host)
  - Creating volumes and attaching them to hosts

Follow the on-screen directions to complete setting up your system. Once you complete the preboarding and onboarding steps you will be taken to the system "Dashboard" on page 42. Here is where you begin to use the SMC to monitor, manage, and provision the storage system.

IMPORTANT Performance of the system will be reduced until initialization of all disk groups is complete.

Using the interface

This section specifies web-browser requirements, describes the user interface, and provides tips for using it.

Web browser requirements and setup

Supported browser versions:

- Apple Safari 11 and newer (Mac)
- Google Chrome 70 and newer
- Microsoft Edge 80 and newer (Chromium-based versions)
- Mozilla Firefox 68 and newer

For best results, use the following guidelines:

- The recommended resolution for the browser's page display area is 1360 x 768 pixels.
- To see the help window, enable pop-up windows.
- 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 SMC is configured to use HTTPS, ensure that your browser is set to use TLS 1.2 or TLS 1.3.
  - Verify that the browser is set to allow cookies, at least for the IP addresses of the storage system network ports.

NOTE By default, your system is loaded with self-signed certificates. We recommend that you:

1. Generate a certificate signing request.
2. Have the certificate signed.
3. Upload the certificate(s) of the chain of trust Certification Authorities (CAs) if necessary.
4. Upload the certificate to the storage system.

Areas of the interface

The main areas of the SMC 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, select the related link in the table.

Selecting an option on the menu pane expands a dropdown list of menu choices. Selecting a menu option displays applicable content in the 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 selecting the slide-over arrows to view more information about the system status and make applicable changes to system settings and configuration. You can select the information icon to view content that defines or explains more information about a feature/option. For more information about the icons used in the interface, see "Icons in the interface" below.

Table 1 Areas of the SMC interface

Banner:	Product name	? Help	"Setting the date and time" on page 57	"User settings" on page 54	**Log Out**
Menu pane:	"Dashboard" on page 42	**Management pane**			
	"Provisioning" on page 45				
	"Settings" on page 51				
	"Maintenance" on page 62				

Icons in the interface

The table below displays a list of the most common icons found in the SMC.

Table 2 Icons in the interface

Icon	Name	Use
⊗	Abort/Cancel	Aborts or cancels an operation.
↲	Apply	Applies an edited operation or selection.
X	Cancel	Cancels an edited operation or selection.
!	Critical	Indicates that the item's health is critical, or that an event has Critical severity.
▼	Collapse	Collapses a table row to hide information about an object.
⌂	Dashboard	Indicates that the Dashboard menu is selected.
⊖	Degraded	Indicates that the item's health or the System's overall health is degraded.
🗑	Delete	Lets you delete a value or object.

Table 2 Icons in the interface (continued)

Icon	Name	Use
◎	Disk	Indicates an operation was performed on a disk.
◎	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.
▶	Expand	Expands a table row to provide more detail about an object.
⇧	Export/Upload	Lets you export or upload a file.
★	Favorite	Indicates that the graph selected is a favorite and will display in the dashboard's compact view.
✓	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.
☰	Provisioning	Indicates that the Provisioning menu is selected.
🕒	Schedule	Indicates that a specified task will take place at defined times.
🔒	Secured	Indicates that the system is secured using FDE.
⚙	Settings	Indicates that the Settings menu is selected.
› or ‹	Slide-over arrows	Opens or closes a panel that contains detailed information about an object.
🔓	Unsecured	Indicates that the system is not secured using FDE.
☰	Volume	Identifies the primary volume.
!	Warning	Indicates that an event has Warning severity.

Tips for using the SMC

- Do not use the browser's Back, Forward, Reload, or Refresh buttons. The SMC 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 SMC 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 SMC session.
- As you set options in panels, the SMC 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.
- Select the › icon to expand a panel and view additional details or perform actions. Select the ‹ icon to collapse a panel and view summary information.

- Select the ▶ icon to expand a table row or container and see additional details or perform actions. Select the ▼ icon to collapse a table row or container and hide detailed information.
- Select the ⓘ icon to open the information window to learn more about an option. Select the icon again to close the information window.
- Select the ✎ icon to edit content within a text box or table.
- In the **Hardware** panel (**Maintenance** > **Hardware**), select a component such as an enclosure or disk to display information about that component.
- If your session is inactive for too long, you will be signed out automatically. This timer resets after each action you perform. One minute before automatic sign-out you will be prompted to continue using the SMC.

Tips for using tables

- Select the ▶ icon to expand a table and see additional details or perform actions. Select the ▼ icon 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, select the arrow icon in the column heading to reorder items from low to high. Select 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 dropdown 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.
- To scroll through a table, click within the table and scroll.

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.

For more information see "Downloading historical disk-performance statistics" on page 81.

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. English uses 1 byte per character.

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. When entering storage-space sizes only, either base-2 or base-10 units can be specified.

TIP Management interfaces enable users to set the base for entry and display of storage-space:

- In the SMC, select **Settings** > **Users** > **Local** > **Add New User** > **Base Preference** to set the base preference for a user.
- In the CLI, enter the `create user` or `set user` command with the `base` parameter to set the base preference for a user.

Not all versions of the SMC provide a setting for this purpose. In such cases, use the CLI instead.

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)	$1{,}000^6$

Signing in and signing out

Multiple users can be signed in to each controller simultaneously.

For each active SMC session, an identifier is stored in the browser. All instances of Firefox, Chrome, Edge, and Safari share the same SMC session. They do not run multiple simultaneous sessions.

NOTE If the initial user has not been created, see "Configuring and provisioning a new storage system" on page 10 for directions on signing in to the system for the first time. Otherwise, see the following procedure.

To sign in:

1. In the web browser address field, type `https://<controller-IP-address>`, then 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 SMC sign-in page displays. If the sign-in page does not display, verify that you have entered the correct IP address.

   NOTE HTTPS is enabled by default. To enable HTTP, see "Enabling or disabling system-management services" on page 53 or see the `set protocols` CLI command.

2. On the sign-in page, enter the username and password of an authorized user.

   NOTE A local username is case sensitive and can have a maximum of 29 bytes. The name cannot already exist in the system, include spaces, or include the following: " , < \ :

   NOTE A local password is case sensitive and can have 8 to 32 characters. If the password contains only printable ASCII characters, then it must contain at least one uppercase character, one lowercase character, one numeric character, and one non-alphanumeric character. A password can include printable UTF-8 characters except for the following: a space or " ' , < >

3. Select **Log In**. If the user authentication fails, a message indicates the system is unable to authenticate the login. If the system is available, the **Dashboard** displays. Otherwise, a message indicates that the system is unavailable.

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

2 System concepts

This section provides overviews of system features and concepts.

Linear storage

Linear storage is a method of mapping logical storage 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. 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 pools to configure storage.

IMPORTANT Storage enclosures support Seagate-branded drives only.

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.

All disks in a disk group must be the same type (SATA, 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 6TB disk and one 8TB disk is equivalent to a disk group composed of two 6TB disks. To maximize disk usage, use disks of similar size.

For more information, see:

- "ADAPT" on page 20
- "Disk sector format" on the facing page
- "Linear disk group configuration" on the facing page
- "Disk-group expansion" on page 23
- "Disk-group scrub" on page 23

Linear disk groups

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

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. To dissolve 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.

NOTE If using the **Highest Capacity**, **Highest Performance**, or **Highest Reliability** predefined configuration, you must dissolve the whole configuration in order to delete the disk group(s) and volume(s).

The RAID levels for linear disk groups created through the SMC 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's 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.

Disk sector format

The system supports 512-byte native sector size disks, 512-byte emulated sector size disks, 4K native 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.
- 4K: All disks use 4096-byte native sector size. Each logical block and physical block is 4096 bytes.
- 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.

Linear disk group configuration

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

During onboarding, you are allowed to select one of the following options for configuring storage. Each option is described—replete with pertinent technical details—within its SMC pane.

For enclosures equipped with SAS drives, select one of the following options:

- **Highest Capacity** (default): This predefined ADAPT configuration uses interleaved volumes and automates the creation of all necessary disk groups and volumes.
- **Highest Sequential Performance**: This predefined ADAPT configuration uses non-interleaved volumes and automates the creation of all necessary disk groups and volumes.
- **Highest Reliability**: This predefined ADAPT configuration uses non-interleaved volumes and automates the creation of all necessary disk groups and volumes.
- **Manual** (advanced): This option enables interactive configuration of storage.

For enclosures equipped with variable-capacity SATA drives, select one of the following options:

- SATA Configuration 1: This predefined ADAPT configuration uses interleaved volumes and automates the creation of all necessary disk groups and volumes.
- SATA Configuration 2: This predefined ADAPT configuration uses non-interleaved volumes and automates the creation of all necessary disk groups and volumes.

See also "Configuring and provisioning a new storage system" on page 10. If the storage system prerequisites are not met for **Highest Capacity** or **Highest Sequential Performance**, the storage system configuration is set to **Manual**. If you want to use a different storage configuration, select **Maintenance** > **Reconfigure Storage Settings**, which allows you to delete the configuration and start over.

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 select **Add Disk Group**.

The RAID levels for linear disk groups created through the SMC 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's disk configuration supports it. If RAID 10 is specified, the disk group has a minimum of two subgroups. Additionally, for enclosures equipped with SAS drives, 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.

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" on page 20.

TIP Choosing the right RAID level for your application improves performance.

These disk group creation capabilities apply to enclosures equipped with SAS or fixed-capacity SATA drives. In the SMC 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).

Table 4 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 5 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 sub-groups	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	20 12	Distributed erasure coding with dual disk failure protection using an 16+3, 16+2, 8+3, and 8+2 stripe width	Very fast rebuilds, no spare disks (built-in spare capacity), large storage pools, enhanced parity options, simplified initial deployment and expansion	Requires a minimum of 20 disks for 16+3 or 16+2. Requires a minimum of 12 disks for 8+3 or 8+2. This provides 2 disks of spare capacity.

IMPORTANT RAID 5 and RAID 6 disk groups, which have parity disks, should be created using the power of 2 best practice to align properly. Failure to follow this method can result in significant degradation of sequential write performance. The power of 2 best practice pertains to HDD disk groups.

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 6 Linear disk group expansion by RAID level

RAID level	Expansion capability	Maximum disks
NRAID	Cannot expand.	1
0, 5, 6	You can add 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.	106

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 an ADAPT disk group must be the same type (SATA, SSD, enterprise SAS, or midline SAS), but can have different capacities, provided the range of difference does not exceed a factor of two. For example, mixing a 600GB disk and a 1.2TB disk is acceptable; but mixing a 6TB disk and a 16TB 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 We recommend not mixing 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. 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.

ADAPT interleaved volumes

ADAPT can take advantage of interleaving to provide improved performance. Interleaving is a method of disk storage that puts sequential information into nonsequential disk sectors, which results in faster read times.

Interleaved volumes are used with applicable ADAPT disk groups:

- For enclosures equipped with SAS drives, **Highest Capacity** is configured during initial storage system provisioning through the SMC. This option is available only if:
  - The storage system is fully populated with identical working disks.
  - The storage system is not configured (no existing volumes or disk groups).

  For information about the onboarding and provisioning process, see "Configuring and provisioning a new storage system" on page 10.
- For enclosures equipped with variable-capacity SATA drives, SATA Configuration 1 is applied during initial storage system provisioning through the SMC.

  For information about the onboarding and provisioning process, see "Configuring and provisioning a new storage system" on page 10.
- The `add disk-group` command is used through the CLI with an `interleaved-volume-count` parameter setting greater than 0.

  NOTE When you use this method, the command creates the number of volumes specified in the `interleaved-volume-count` parameter. For more information, see the *CLI Reference Guide*.

With each method, the system automatically creates the volumes sized for best performance.

You can check interleaved volume settings in the SMC in two places.

- In the **Provisioning** > **Volumes** panel, the **Overview** tab of the volumes slide-over shows the Interleaved setting, which is `True` when interleaved volumes are in use, or `False` if they are not.
- When you access a disk group slide-over panel under **Maintenance** > **Storage**, the Interleaved Volume Count setting displays on the **Overview** tab when interleaved volumes are in use. The value shown is the number of volumes in the interleaved disk group. If the disk group does not use interleaved volumes, the Interleaved Volume Count field does not display.

The following limitations apply to ADAPT interleaved volumes:

- An interleaved ADAPT disk-group cannot be expanded.
- A volume that is part of an interleaved ADAPT disk-group cannot be expanded.
- ADAPT interleaved volumes cannot be deleted individually. Instead, you must delete the disk group and create new volumes and disk groups with desired characteristics.

With ADAPT interleaved volumes, you can change volume names and mapping through the SMC just as with other types of volumes.

SPEAR

SPEAR is a data-protection feature you can enable for disk groups that ensures data integrity by verifying consistency of the stripe with each read access. Stripe Parity Error Assessment on Read (SPEAR) checks parity of blocks in stripes before returning data to the host.

You can enable and configure SPEAR for a disk group from the **Maintenance** > **Storage** panel. When creating a new disk group, set up SPEAR on the **Add Disk Group** dialog box. For existing disk groups, change SPEAR settings on the disk group slide-over panel.

Why Use SPEAR?

SPEAR checks parity on every read access, so verification occurs during host read operations as well as during some host writes. Additionally, internal utilities such as reconstruction, rebalance, expand, and scrub all do reads. Any data read from disk is validated by SPEAR.

When an inconsistency is found, the controller fixes the error if possible, keeping data integrity as the priority. Correction is possible only if there are no drive errors in the stripe and if all data blocks are consistent with at least one of the parity blocks.

SPEAR detects problems with stale write data and silent corruption in parity blocks that foundational T10 Protection Information (PI) technology does not. However, the additional data integrity can affect performance.

Stale write data can occur if a write request completes successfully without actually updating the data on the disk. With SPEAR enabled, these "dropped writes" are detected on the next read and the data can be recovered from parity.

Silent corruption of the parity blocks can occur with PI technology. Although PI can be used to rebuild, it isn't valid for the data in the parity block itself. Therefore, checking is turned off when writing data, which is how silent corruption can occur. With SPEAR enabled, any corrupted parity blocks are caught and can be repaired on the next read.

Because extra reads are required to check stripe parity with SPEAR, system throughput and I/O performance can be lower. The performance impact is dependent on host access patterns. Host requests that access large amounts of sequential data see little impact, while requests for smaller amounts of random data typically have higher impact.

System requirements

To enable SPEAR, the disk group must meet the following minimum requirements:

- RAID 6 or ADAPT
- 4KB drive sectors
- Linear disk groups only (virtual disk groups are not supported)
- PI type 0 and 1 on controller host ports

Types of errors

SPEAR can encounter two types of problems when attempting a read:

- `Unavailable`: The algorithm is not able to check stripe integrity because two or more disks (also known as blocks or symbols) in the stripe are missing. Disks can be unavailable due to disk failure or other errors.
- `Error`: Enough drives are read but a parity check fails. In this case, the stripe data is not consistent with one or both of the parity blocks, which indicates the data could be corrupt.

When an error condition is encountered, requested read information might still be returned, based on the current mode of operation (see below).

Modes of operations

You can operate SPEAR at three different levels of protection:

- Basic: Returns requested user data unless SPEAR error is detected. In the case of SPEAR unavailable, unverified data is returned.

- Strict: Returns requested user data unless SPEAR error or SPEAR unavailable is detected.
- Trust: Returns requested user data even if SPEAR error or SPEAR unavailable is detected. Note that in such cases, returned data could be corrupt.

NOTE You can select Trust mode for a disk group only if SPEAR is set to Basic. When Trust mode is set for a disk group, you can transition to Basic or you can disable SPEAR for that disk group, but you cannot change to Strict mode.

You can change SPEAR settings on the slide-over panel for a disk group (**Maintenance** > **Storage**).

Disk group utilities

This section provides information about disk group utilities.

Disk-group expansion

When expanding a disk group, all disks in the group must be the same type (SATA, SSD, enterprise SAS, or midline SAS). 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" on page 16.

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.

NOTE If using the **Highest Capacity**, **Highest Sequential Performance**, or **Highest Reliability** configurations, you cannot expand disk groups because all available disks were used by the configuration algorithm.

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, select **Expand Disk Group**, and follow the on-screen directions.

Expanding non-ADAPT linear disk groups

Before expanding non-ADAPT disk groups, back up the disk group's data 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, select **Expand Disk Group**, and follow the on-screen directions.

To cancel disk group expansion, select the ⊗ icon.

Disk-group scrub

The disk-group scrub utility analyzes a disk group to find and fix disk errors. It acts on all disks in the disk group but not dedicated spares for a linear disk group or leftover disks.

The disk-group scrub utility analyzes specified disk groups to find and fix errors.

TIP The disk-group scrub utility can find media errors for any protection level. 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 will:

- Check redundancy data (parity) and correct it for protection levels 5, 6, and ADAPT
- Find, but not fix, mirror mismatches for protection levels 1 and 10

  The system reads both copies of mirror data to find any mismatches.
- Find and fix 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).

The disk-group scrub task is almost continually active, but has very little impact on I/O activity. You can use a disk group while it is being scrubbed. While scrub is running, you can monitor progress and cancel if necessary. When scrub is complete, event 207 is logged. Follow recommended Error or Warning conditions described for event 207 in the Event Descriptions Reference Guide.

To run the scrub utility, see "Scrubbing a disk group" on page 64.

Autonomous drive regeneration

Autonomous Drive Regeneration (ADR) repairs a spinning disk that has become unusable due to a single head failure. ADR is supported for ADAPT disk groups only.

CAUTION ADR will not run on a storage system that uses FDE drives and is in a secured state.

Rather than replacing the disk that has incurred a head failure, this utility enables removal of the bad head to provide a remanufactured disk—with contiguous LBA range—albeit of less capacity than the original disk.

TIP ADAPT does not require the disks in a disk group to be of the same capacity, so the remanufactured disk is seamlessly added back to the disk group once the process completes.

An enable/disable checkbox-toggle controls the use of Remanufacture (**Settings** > **System** > **Properties** > **Disk Properties**). The progress of a remanufacture process is reported in **Dashboard** > Activity.

The remanufacture utility will:

- Function as a system-wide setting addressing spinning disks included in ADAPT disk-groups.
- Operate on a disk that has been removed from an ADAPT disk-group due to head failure.
- Process only one degraded disk with a failed head at a time.
- Operate automatically based on the status of the enable/disable checkbox.

  If enabled, the controller automatically takes the disk offline and initiates the two-phase remanufacture process. Phase 1 logically removes the damaged head and is short in duration. Phase 2 reformats the newly sized disk and can take many hours.

  During remanufacture, the controller uses spare capacity in the ADAPT pool to keep the disk-group fault-tolerant. Upon completion, the controller brings the remanufactured disk online, reestablishing it in the same ADAPT disk-group to which it previously belonged. An ADAPT rebalance operation will start immediately, to move data onto the newly added disk—balancing data between all disks in the pool—and restoring spare capacity.

NOTE If the owning disk-group is removed before completion of remanufacture on a degraded disk from that disk-group, the disk will return to Available state.

When remanufacture is started, event 630 is logged. When remanufacture is complete, event 631 is logged. Follow recommended `Error` or `Warning` conditions described for these events in the Event Descriptions Reference Guide.

SSDs

ⓘ **IMPORTANT** Seagate Exos CORVAULT storage systems do not use SSDs. The discussion provided herein is for informational purposes only.

The use of solid-state drives (SSDs) may 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 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.

For more information, see:

- "Gauging the percentage of life remaining for SSDs" below
- "SSD management" below

Gauging the percentage of life remaining for SSDs

An SSD can be written and erased a limited number of times. Through the SSD Life Remaining 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 SMC, select the SSD. See the SSD Life Remaining 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.

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.

Write amplification

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.

`TRIM` and `UNMAP` commands

A command (known as `TRIM` in the ATA command set and `UNMAP` 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 disk is powered off is affected by Program and Erase (PE) cycles and the temperature of the disk 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.

Because data is characterized every 5 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.

Spares

Spare disks are unused disks in your system that automatically replace a failed disk, restoring fault tolerance to disk groups in the system. Only the Manual configuration allows for creating spares within the SMC. Designate spares from the **Maintenance** > **Storage** panel. Designate spares 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" on page 20.

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 speed and type (SATA, SSD, enterprise SAS, or midline SAS). It is not advisable to mix 10k and 15k disks in a single disk group. 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.

TIP 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.

In the SMC you can designate a maximum of 64 global spares. If a disk in any fault-tolerant disk group fails, a global spare (which must be the same size or larger and the same type as the failed disk) is automatically used to reconstruct the disk group (RAID 1, 5, 6, 10). At least one disk group must exist before you can add a global spare. A spare must have sufficient capacity to replace the smallest disk in an existing disk group.

The disk group will remain in critical status until the data, parity, or mirror data is completely written to the spare, at which time the disk group will return to fault-tolerant status.

Disk groups support a mix of 512n and 512e disks. However, for consistent and predictable performance, do not mix disks of different rotational speed or sector size types (512n, 512e). If a global spare has a different sector format than the disks in a disk group, an event will appear when the system selects the spare after a disk in the disk group fails. For more information about disk groups, see "Disk groups" on page 16.

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 linear pools, which can only have one disk group per pool, volumes are also added to the pool, which contains the volume data.

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 configure disks into disk groups. For information about how provisioning disks works, see "Adding a disk group" on page 63.

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.

Volumes

A volume is a logical subdivision of a 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 about attaching volumes to hosts, see "Attaching volumes to hosts" on page 30.

For linear pools, which can only have one disk group per pool, volumes are also added to the pool, which contains the volume data.

Linear volumes make use of a method of storing user data in sequential fully allocated physical blocks. 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" below
- "Cache optimization mode" below
- "Optimizing read-ahead caching " on the facing page

Using write-back or write-through caching

⚠ **CAUTION** Only disable write-back caching if you fully understand how the host operating system, application, and adapter move data. If used incorrectly, you might hinder system performance.

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 actually written to the disk. Write-back cache is saved to non-volatile storage in the event of a power loss. Write-back cache mirrors all of the data from one controller module cache to the other in the event of a controller fault, and the remaining controller completes the write operation to the disks. When modifying a volume you can change its write-back cache setting. Write-back cache improves the performance of write operations and the throughput of the controller.

When write-back cache is disabled, write-through becomes the cache-writing strategy. Using write-through cache, 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. However, write-through cache does not mirror the write data because the data is written to the disk before posting command completion and mirroring is not required. You can set conditions that cause the controller to change from write-back caching to write-through caching. For more information, see "Setting system cache properties" on page 59.

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

You can enable and disable the write-back cache for each volume. By default, volume write-back cache is enabled. Because controller cache is backed by supercapacitor technology, if the system loses power, data is not lost. For most applications, this is the preferred setting.

TIP 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 integrity issues or data 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's 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

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

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 RAID-1 disk groups internally as if they have a stripe size of 512KB, even though they are not striped.
- Specific size options let you select an amount of data for all accesses. Options include 512KB, 1MB, 2MB, 4MB, 8MB, 16MB, 32MB.
- 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.

Hedged reads

The hedged reads feature is a method for read requests that takes advantage of RAID data to speed data retrieval. When a read is requested, the system waits a predefined period for the read to return the data. If the threshold passes before the data is returned, data reconstruction takes place instead. In this way, the system limits latency from slowly responding HDDs and returns data faster.

The following limitations apply to hedged reads:

- Limited to HDDs. Does not function for SSDs.
- Applies only to RAID levels 1, 5, 6, 10, and ADAPT.
- Used for host reads only. Does not include utility reads or reads associated with partial stripe writes.

NOTE Hedged reads functionality operates by default when the conditions are met. There are no user settings or controls. For assistance with hedged reads, contact technical support.

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 an HBA) in a server.

For ease of management, you can group from 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 group 1 to 128 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 `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 SMC, 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.

Host ports

Controller enclosures support SAS host interface protocol. SAS controller host-port settings are not configurable in the SMC.

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" on page 47.

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).

ⓘ **IMPORTANT** 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

Seagate Exos enclosures support dual-controller configurations only. If a partner controller fails, the storage system will fail over and run on a single controller until the redundancy is restored. A controller module must be installed in each IOM slot to ensure sufficient airflow through the enclosure during operation.

Reconstruction and copyback

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" on page 26).

For descriptions of LED states, such as for disk failure and reconstruction, see the *Hardware Installation and Maintenance Guide*.

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. When the system detects the new disk, it initiates a copyback operation, which copies all data to the new disk from the spare disk that replaced the failed disk. When the copyback operation is complete, the spare disk is freed so that it can be used for a subsequent disk failure.

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 than with traditional RAID protection levels. 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.

Preemptive reconstruction

Preemptive reconstruction provides a way to preemptively replace a potentially defective disk in a RAID 1, 5, 6, or 10 disk group. This command is not supported for a RAID-0 disk group. For information, see the topic about the `force disk–degraded` command in the CLI Reference Guide.

Updating firmware

Controller modules and disk modules contain firmware. Users must have a `manage` role to update the disk or system firmware. For information on these options, see:

- "Updating system firmware" on page 66
- "Updating disk firmware" on page 68

For more information, see "Best practices for updating firmware" on page 68.

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 stored for later retrieval before any log 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, IP address and the log-file type (region) that needs to be transferred. To activate pull mode, deselect 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 are 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 a log's data in push or pull mode, the log's 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

Typical log data that can be written to a compressed file include:

- 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 loading firmware. Do not perform more than one log saving operation at a time. Do not try to perform a firmware update operation 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" on page 78.

LDAP

You can configure the storage system to use external Lightweight Directory Access Protocol (LDAP) services provided from Windows Server 2016, 2019, or 2022 Active Directory for user authentication and authorization.

Feature overview

There are two sources of user credentials for the storage system. The primary source is local users created by using the options in the **Local Users** panel of the SMC (**Settings** > **Users** > **Local**) or by using the `create user` CLI command. For more information on this command, see the CLI documentation. For more information on adding local users with the SMC, see "Managing local users" on page 55. Though local users can be standard or SNMPv3 users, the LDAP feature supports only standard users.

The secondary source for user credentials is a Windows 2016, 2019, or 2022 Active Directory LDAP server, as illustrated below. Users logging in using their LDAP credentials must authenticate using these credentials and be members of a group that is authorized to access the storage system. The group will exist on the LDAP server and will be listed under the `memberOf` attribute for the user account. The same group name must also exist in the storage system and be created by using the LDAP Users panel of the SMC (**Settings** > **Users** > **LDAP**) or the `create user–group` CLI command. Users logging in by this method are not explicitly registered or stored in the storage system; their login, logout, and activity is recorded in an audit log stored in each controller module. For more information about audit logs, see "Audit logs" on page 41.

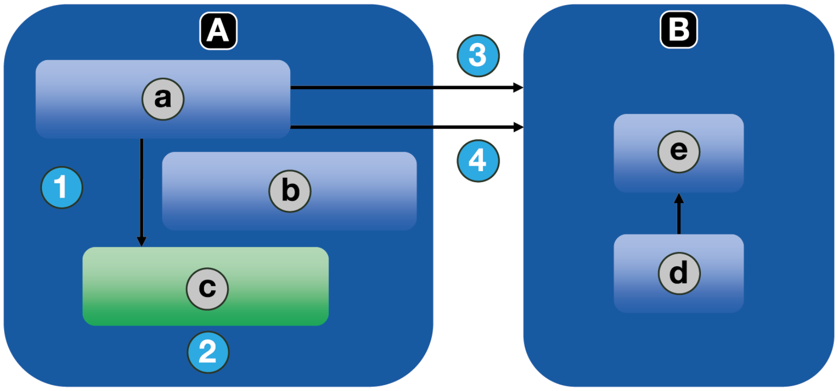

Device		Element		Process	
A	Storage system	a	Session Manager	1	Create session with userid, inherit role from group
B	LDAP server	b	Group: adminusers Role: manage	2	Username will only exist as part of session
		c	Session User: joe Role: manage	3	Authenticate: user
		d	User: joe	4	Check group membership
		e	Group: adminusers		

Figure 1 LDAP overview

The system supports a maximum of five user groups to allow different permissions and/or user preference options. User group permissions are defined by assigning roles, as for local users. User group preference options that can be set in the SMC include the temperature scale, timeout, and storage base size. User group preference options that can be set only in the CLI include the precision and units. User groups can be created whether the LDAP feature is enabled or disabled, but have no purpose if LDAP is not enabled.

Individual user preferences are not saved in the storage system for LDAP authenticated users. Any settings made to the login session are not retained after the session terminates. If the user wants to retain any preferences for the session, these must be saved as part of the user group. Any changes made to a user group will affect all members of that group.

LDAP users with a `manage` role can create, modify, and delete both local users and user groups. LDAP users with a `standard` role can change settings for the current user group except for the user type and role. LDAP users with a `standard` role also cannot change the settings of other user groups.

The username/password entered will be authenticated with local users within the system first. If local authentication fails and LDAP is enabled, the username will be checked against the LDAP server(s).

Protocols and services

Before enabling the LDAP feature, insecure protocols and services (Telnet, HTTP, FTP, and debug) must be disabled.

When the LDAP feature is enabled, only secure protocols and services (SSH, HTTPS, SFTP) can be enabled. The LDAP feature must be disabled before insecure protocols can be re-enabled.

HTTPS, SSH, and SFTP are the only interfaces that support LDAP. Attempting to use an LDAP login through any other interface will fail.

LDAP server/client details

The LDAP server must be an Active Directory server running at a function level of 2016 or later (Windows 2016, Windows 2019, or Windows 2022), and LDAP must be enabled on Windows Server Active Directory before it can be enabled and successfully used by the storage system. The server must allow basic authentication using an LDAP over SSL (LDAPS) interface; that is, a TLS 1.2 or TLS 1.3 connection. The SMC LDAP feature is not intended for use with generic LDAP on Linux.

NOTE As an older protocol, LDAP supports only simple authentication without data encryption. As a newer protocol, LDAPS supports LDAP data encryption and SSL or TLS for the connection handshake.

The client storage system allows one primary server and port and an alternate server and port to be configured. At login, the storage system will only connect over TLS. If the storage system cannot connect to the primary server it will automatically try the alternate server. The storage system will only connect to a single Active Directory forest.

The client will look at the common name (CN) for the LDAP group's distinguished name (DN). The group can be part of any organizational unit (OU) or Active Directory forest as long as the CN value matches the client's group name.

For example, assume domain `bigco2.com.local` includes OU `colo`, in which user `alice` is a member of group `ArrayAdmins` in the same OU. The group's DN is: `cn=ArrayAdmins,ou=colo,dc=bigco2,dc=com,dc=local`

When the LDAP client performs a search on the server, it will query the `UserObject` that represents user `alice`. The client will limit the response to a maximum of 100 groups to be read from the server. The first group found that matches a group created on the storage system will be used to authenticate user `alice`. The client will timeout if it has not received a response in 20 seconds.

In the above example, the user group `ArrayAdmins` has been created on the storage system. When the user `alice` attempts to log in to the storage system either through the SMC or the CLI, the group from Active Directory matches the storage system user group and `alice` is granted access.

It is recommended that:

- A user should only be a member of one group that exists in the storage system. A user that is a member of more than one LDAP group in the storage system could have permission or configuration parameter inconsistencies.
- The LDAP user be in no more than 100 LDAP groups.

The following example shows the data to enter in the LDAP Configuration panel to configure a storage system to accomplish the above.

1. Configure the storage system to connect to the primary LDAP server and an alternate LDAP server. IP addresses or Fully Qualified Domain Name (FQDN) may be used. The primary connection is configured at 10.235.217.52 using standard TLS port 636. The alternate connection is configured at 10.235.217.51 using the same port. If the primary connection fails, the system will try the alternate connection. If the alternate connection fails, authentication will fail. The user search base defines the domain and OU.
   a. Access the LDAP Settings section via **Settings** > **Users** > **LDAP**.
   b. Select the **Enable LDAP** check box.
   c. In the User Search Base field, enter **`ou=colo,dc=bigco2,dc=com,dc=local`**.
   d. In the Server field, enter **`10.235.217.52`**.
   e. In the Port field, enter **636**.
   f. In the Alternate Server field, enter **`10.235.217.51`**.
   g. In the Alternate Port field, enter **636**.
   h. Select **Set LDAP**.

2. Create an LDAP user group named `ArrayAdmins` (matching the group name on the LDAP server) with the Standard role and access to the SMC and CLI interfaces.
   a. Select **Add New User Group**.
   b. In the **User Group Name** field, enter **`ArrayAdmins`**.
   c. Select **WBI** and **CLI** to define the interfaces.
   d. Select **Standard** and **Monitor** to define the roles.
   e. Select the temperature preference and timeout options.
   f. Select **Create User Group**. When user `alice` attempts an SSH login to the storage system, the system connects to the configured LDAP server using the supplied credentials to perform authentication.

There are two login formats that the storage system allows when connecting to an Active Directory LDAP server. When using SSH, two backslashes may be required for certain clients, such as the OpenSSH client.

- Email-address format. For example:

  ```
  ssh alice@bigoc2.com.local@10.235.212.161
  ```

- Domain\username format. For example:

  ```
  ssh bigco2\\alice@10.235.212.161
  ```

Using the domain\username format has this restriction: the username can contain no more than 20 characters to be backward-compatible with Windows clients before Windows 2000. For more information about restrictions for these attributes, see Microsoft Active Directory documentation.

NOTE By default when creating a new user object in Windows Server 2016 or 2019, both the `sAMAccountName` and `userPrincipalName` attributes are populated.

Recovery

If the LDAP server becomes permanently unavailable or no users exist in the LDAP database and local user account passwords are forgotten or compromised, physical access to a controller module serial port will be required. If this occurs, contact technical support for assistance.

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.

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.

⚠ **CAUTION** ADR will not run on a storage system that uses FDE drives and is in a secured state.

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 disk's encrypted data is inaccessible.

ⓘ **IMPORTANT** 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" below.

Rescanning disks

A rescan (**Maintenance** > **Hardware** > **Actions**) forces a rediscovery of disks.

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. Selecting 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 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's health becomes `Degraded` and its usage value becomes `LEFTOVR`.
- The disk is automatically excluded from the disk group, causing the disk group's health to become `Degraded` or `Fault`, depending on the RAID level.
- The disk's fault LED 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 disk's metadata. Clearing the metadata will delete all data on the disk and change the disk's health to `OK` and its usage value to `AVAIL`. The disk may become available for use in a new disk group.

TIP 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.

NOTE 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 a disk's metadata has been cleared, verify that the disk's health is `OK`.
- When rebuilding a disk group, do not use an unhealthy disk from which metadata has been cleared.

Data protection with a single controller

The system can operate with a single controller if its partner has gone offline or has been removed. Because single-controller operation is not a redundant configuration, this topic presents some considerations concerning data protection.

The default caching mode when a system is operating with a single controller for a volume is write back, as opposed to write through. In write-back mode, the host is notified that the controller has received the write when the data is present in the controller cache. In write-through mode, the host is notified that the controller has received the write when the data is written to disk. Therefore, in write-back mode, data is held in the controller cache until it is written to disk.

If the controller fails while in write-back mode, unwritten cache data likely exists. The same is true if the controller enclosure 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.

To avoid the possibility of data loss in case the controller fails, you can change the caching mode of a volume to write through. While this will cause a performance degradation, this configuration guards against data loss. While write-back mode is much faster, this mode is not guaranteed against data loss in the case of a controller failure. If data protection is more important, use write-through caching. If performance is more important, use write-back caching.

For more information about volume cache options, see "Volume cache options" on page 28. To edit volume cache options, select **Provisioning** > **Volumes** and view the volume's slide-over panel.

For more information about changing system cache settings, see "Setting system cache properties" on page 59.

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" on page 61.

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" on page 14.

Table 7 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, select 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" on page 42.

Audit logs

User login, logout, and actions through all interfaces for both local and **LDAP** users will be recorded in an audit log that is accessible from **Maintenance** > **Support** > **Audit Logs**. 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 SMC, CLI, and SFTP/FTP 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's 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.

3 Dashboard

Use the system dashboard to monitor the system and see an overview of system status, including:

- Health and performance alerts in the "Alerts panel" below
- Trends in capacity usage in the "Capacity panel" on the facing page
- System activity in the "Activity panel" on page 44

Each panel 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.

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 "Icons in the interface" on page 12.

Compact view

The compact view on the dashboard provides a snapshot of the system's overall health, including a count of health alerts, information alerts, and unacknowledged alerts. Select **System Health | Health Alerts** to view a scrollable list of unresolved health alerts that are affecting system health. Select **Information Alerts** to view a scrollable list of unresolved informational alerts that notify you of actions to take to improve system performance. Select **Alerts to Acknowledge** to view a scrollable list of all resolved and unresolved health and information alerts that need acknowledgment. Once acknowledged and resolved, alerts are removed from the list.

Select **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 upon the type of alerts you want to view (**Health Alerts**, **Information Alerts**, **Alerts to Acknowledge**, or **History**). Select 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**)
- Whether the alert has been acknowledged (**Health Alerts**, **Information Alerts**, **Alerts to Acknowledge**) (**Alerts to Acknowledge**: present if not acknowledged; not present if acknowledged, and checked in other lists)

Select ▶ 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
- A field where you can view or enter a comment about the alert, such as direction to another user

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

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

To enter or edit a comment in the COMMENTS field, enter text in the field and then select ↲ to save or X to cancel. To delete a comment, select 🗑.

Select **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 to end users and support personnel. Historical alerts remain in the **History** table until it reaches its threshold limit of 512. When it does, it begins to wrap.

Capacity panel

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

Compact view

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).

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

Select the › icon to see the expanded view.

Expanded view

The expanded view shows:

- Disk-group allocation (select 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 (select the arrow to expand the row). For each configured volume, the panel shows:
  - A filtering field with dropdown selection.
  - A table providing data for each volume: name, size, percentage of disk group, and attached hosts.

NOTE The system requires metadata space for volume writes. When the metadata value reaches zero, the system cannot perform further volume writes, regardless of the amount of free disk space. Accordingly, the system issues threshold events if metadata usage exceeds certain levels. If this happens, it becomes necessary to move or remove lower priority volumes to prevent stoppage of further volume writes. Once this is done, the metadata metrics should recalibrate to display updated values.

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. A progress bar shows status using green fill and text reporting percent complete.

  NOTE The SMC 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.

Select the › icon 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.

The 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.

The 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.

The 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 dropdown list enables selection of activities to display and monitor.

For tips on using tables, see "Tips for using tables" on page 14.

4 Provisioning

Use the **Provisioning** panel to view, create, and manage volumes and hosts. For more information, see:

- "Working with volumes" below
- "Working with hosts" on page 47

Working with volumes

The **Volumes** panel (**Provisioning** > **Volumes**) provides options to create volumes and to display a list of existing volumes on the system. For information about creating volumes, see "Creating volumes" on the next page.

If storage is not configured, the panel displays a prompt for you to add storage. If 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, select one or more volumes and select an option from the dropdown list:

- "Deleting volumes " on the next page
- "Attaching volumes to hosts" on page 47
- "Detaching volumes from hosts" on page 47
- "Expanding volumes" on page 47

Other actions to perform on an existing volume include:

- "Modifying volumes" on the next page

Select the volume slide-over › to view volume details and to perform more actions on a volume .

- The **Overview** tab lists all of the volume settings. Select the ✎ icon to make changes to these settings. Select **Expand Volume** to expand the size of the volume. Follow the on-screen directions for more details. For more information, see:
  - "Cache optimization mode" on page 28
  - "Optimizing read-ahead caching " on page 29
- The **Attached Hosts** tab displays a table listing the hosts attached to the selected volume and lets you attach the volume to a host or host group. Hover over icons in the list to see more information. Select the ✎ icon 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 selectable option allowing you to detach the host from the volume.

For more information about hosts and initiators, see "Initiators, hosts, and host groups" on page 29.

Creating volumes

Select **Create Volumes** (**Provisioning** > **Volumes**) to open the **Create Volumes**) wizard to add a volume to a disk group. The top portion of the panel displays a color-coded graph for volume usage by disk group. The graph indicates the amount of space on the system that is allocated to volumes; the space occupied by added volumes; the space required for the new volume you are creating; and available space.

NOTE For interleaved disk groups, you must use the CLI to create individual volumes. The **Create Volumes** wizard is not available.

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. Select **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" on page 27
- "Initiators, hosts, and host groups" on page 29

NOTE Volume sizes are aligned to 4.2MB (4MiB) boundaries. When a volume is created or expanded, the resulting size will be decreased to the nearest 4.2MB boundary. For the maximum volume size supported by the system, see "System configuration limits" on page 94.

Modifying volumes

Change the volume settings from the Volumes table (**Provisioning** > **Volumes**) by selecting the volume's slide-over to access the **Overview** panel. Here you can expand volumes, edit the volume name, and select cache setting options.

For more information, see "Volume cache options" on page 28.

⚠ **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.

Deleting volumes

You can delete volumes from the Volumes table (**Provisioning** > **Volumes**):

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

⚠ **CAUTION** Deleting a volume removes its host attachments and schedules and deletes its data.

Select a volume, then select **Delete Volumes** from the dropdown list. Follow the on-screen directions to complete the action. The following rules apply:

- You can select from 1 to 100 items (volumes) to delete.
- Ensure that hosts are not accessing the volumes to be deleted.

Attaching volumes to hosts

Attach volumes to hosts from the:

- Volumes table (**Provisioning** > **Volumes**). Select the volume and select **Attach to Hosts** from the dropdown 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:

- Volumes table (**Provisioning** > **Volumes**). Select a volume from the Volumes table and select **Detach from Hosts** from the dropdown list.
- Attached Hosts panel (**Provisioning** > **Volumes** > 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

NOTE The capability to expand volumes does not apply to **Highest Capacity** or **Highest Performance** configurations; nor does it apply to a volume that is part of an interleaved disk-group.

You can expand the size of a volume from the:

- Volumes table (**Provisioning** > **Volumes**). Select the volume and select **Expand Volumes** from the dropdown list.
- **Overview** panel (**Provisioning** > **Volumes** > 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 volume usage by disk group. The graph indicates the amount of space on the system that is allocated to volumes; the space occupied by added volumes; the space required for the new volume you are creating; and available space.

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

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" on the next page. 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 select an option from the dropdown list:

- "Attaching hosts to volumes" below
- "Detaching hosts from volumes " on the facing page
- "Removing initiators from a host" on the facing page
- "Removing hosts from a host group" on the facing page
- "Adding hosts to a host group" on the facing page
- "Deleting hosts" on the facing page
- "Deleting host groups" on the facing page

Other actions to take on this tab include:

- "Renaming hosts " on page 50
- "Changing a host profile" on page 50
- "Renaming host groups" on page 50
- "Renaming initiators" on page 50

Select the ▶ icon to expand the host row to see initiator details. Select a host or initiator to perform an action from the dropdown list.

Select the host slide-over to view the **Overview** tab, where you can edit the name of the host and nickname of each initiator. Select 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.

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

Select 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 select an option from the dropdown list:

- "Adding initiators to a host" on the facing page
- "Removing initiators from a host" on the facing page

For more information about hosts and initiators, see "Initiators, hosts, and host groups" on page 29.

Creating hosts

Select **Create Hosts** (**Provisioning** > **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.

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

For more information about volumes, see "Volumes" on page 27. For more information about hosts, host groups, and initiators, see "Initiators, hosts, and host groups" on page 29.

Attaching hosts to volumes

Attach hosts to volumes from the Hosts table (**Provisioning** > **Hosts**) by selecting the host and selecting **Attach Volumes** from the dropdown 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 selecting Detach Volumes from the dropdown 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 selecting **Remove From Host** from the dropdown 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 selecting **Remove From Host** from the dropdown 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" below.

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 selecting **Add To Host Group** from the dropdown 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 selecting **Delete Host** from the dropdown 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 selecting **Delete Host Group** from the dropdown 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.

Adding 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 selecting **Add To Existing Host** from the dropdown 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. Select 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). Select ✎ 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). Select ✎ within the Hosts table, then select an option from the Profile dropdown menu.

Renaming host groups

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

Renaming initiators

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

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

5 Settings

Use the **Settings** panel to view and manage system configuration settings, including:

- "Network settings" below
- "User settings" on page 54
- "System settings" on page 57
- "Notification settings" on page 61

Access the panel by selecting the applicable option from the **Settings** menu pane.

Network settings

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

- "Configuring controller network ports" below
- "Configuring DNS settings" on the next page
- "Enabling or disabling system-management services" on page 53
- "Viewing certificate information" on page 54
- "Configuring a proxy server" on page 54

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 (see bullet lists and **IMPORTANT** note provided in the next paragraph). 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

ⓘ **IMPORTANT** 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
- SLAAC IP address: - Not Set -

The Link-Local Address that is also displayed in this grouping of **Auto** address fields is not a DHCPv6 or SLAAC assigned address.

△ **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, see "DNS settings" on page 37.

NOTE DNS settings are automatically applied when using DHCP for IPv4 and Auto for IPv6.

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, select the **Reset Host Name** button for that controller.
- To clear a DNS server or search domain for a controller, select the 🗑 icon for that setting.
- To clear all configured DNS servers and search domains for a controller, select the **Clear DNS** button for that controller.

For more information about the DNS feature, see "DNS settings" on page 37.

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 SMC, CLI, or other interfaces, see "User settings" on the next page.

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

Web and API

- **HTTPS**. Enable the use of HTTPS to provide secure access to the SMC, the web application that is the primary interface for managing the system.
- **HTTP**. Enable the use of HTTP to provide access to the SMC.

NOTE One of the above settings must be enabled to provide access to the SMC.

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.

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.

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. 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 **Certificate Information** panel (**Settings** > **Network** > **Certificates**). By default, the system generates a unique SSL certificate for each controller. For the strongest security, upload a Certificate Signing Request (CSR) signed by a Certificate Authority (CA), along with the certificates of the CA trust chain.

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's certificate cannot be read. This situation occurs most often when a controller is restarting or certificate replacement is still in progress.

You can upload your own certificate generated by creating a CSR through the CLI, sign the CSR using a CA, upload the certificate(s) of the CA trust chain via SFTP/FTP and then upload the signed certificate via SFTP/FTP. Ensure that the CA certificate(s) are uploaded first so they can be used to verify the signed device certificate. 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, or run the CLI command `restart mc full`.

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.

Configuring a proxy server

The **Proxy Server** is a shared configuration option for all storage system features that require a proxy for external network communication. For environments that are behind a firewall and do not have direct access to the Internet, you can configure a proxy server to act as an intermediary between the storage system and the external URL. Only HTTP is supported to the proxy server, though the proxy itself can support HTTPS traffic to the final endpoint.

To configure a proxy server, select **Settings** > **Network** > **Proxy Server**, select the **Enable Proxy** checkbox, and then set the proxy hostname or IP address and port number (if any). If required by the proxy, also enter a username and password. To clear the configured username and password, select **Clear Proxy**. Once the parameters are set, select the **Set Proxy Server** button to complete the task.

To disable a configured proxy server, select **Settings** > **Network** > **Proxy Server** and clear the **Enable Proxy** checkbox.

User settings

The **Users** panel (**Settings** > **Users**) provides options for you to manage local users, LDAP users and user groups, 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" on the facing page
- "Managing LDAP users" on the facing page
- "Managing SNMPv3 users" on page 56

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 the `standard` or `monitor` role can change their own username, password, 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.

Users having the `manage` or `standard` role can access one or more of the following management interfaces: the SMC, CLI, or FTP and SFTP. Monitor users can only access the SMC 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 the `manage` or `standard` role when adding or modifying users. To add new users, select **Add New User**, and to modify users select the ✏ icon.

- **Username**. A username is case sensitive and can have a maximum of 29 bytes. The name cannot already exist in the system, include spaces, start with a hyphen, or include any of the following: " , < \
- **Password**. A password is case sensitive and can have from 8 to 32 characters. If the password contains only printable 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 printable UTF-8 characters except for the following: a space or " ' , < > \
- **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's 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 SMC.
  - **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 SMC 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 LDAP users

The LDAP Configuration panel (**Settings** > **Users** > **LDAP**) provides options for users with the `manage` role to create up to five user groups to allow for different permissions and/or user preference options. User group permissions are defined by assigning roles. User group preference options include selecting interfaces, role, temperature preference, and a timeout setting.

Users logging into the SMC using their LDAP credentials must authenticate using these credentials and be members of a group that is authorized to access the storage system. The username and password entered will be authenticated with local users within the system first. If local authentication fails, the username will be checked against the LDAP server(s).

Individual user preferences are not saved in the storage system for LDAP authenticated users. Any settings made to the login session are not retained after the session terminates. If the user wants to retain any preferences for the session, these must be saved as part of the user group. Any changes made to a user group will affect all members of that group.

To enable LDAP, you must select the **Enable LDAP** checkbox and enter the user search base, server address, and port. If the port is left blank it will default to 636. For more information about these options, see "LDAP" on page 34.

LDAP user group options

As a user with the `manage` role, you can modify or delete any user group. As a user with only a `standard` or `monitor` role, you can change settings for the current user group with the exception of user roles and interfaces. You also cannot change the settings of other user groups.

- **User Group Name**. A user group name is case sensitive and can have a maximum of 29 bytes. The name cannot already exist in the system or include any of the following: " , < \
- **Interfaces**. Select one or more of the following interfaces:
  - **WBI**. Enables access to the SMC.
  - **CLI**. Enables access to the command-line interface.
  - **FTP**. Enables access to the SFTP interface or the FTP interface, which can be used instead of the SMC 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.
  - **Standard**. Enables the user to change system settings except for: creating or deleting local users, modifying user role and interfaces, configuring LDAP, performing write operations through FTP or SFTP, performing file uploads from the SMC, or using the `restore defaults` CLI command.
  - **Monitor**. Enables the user to view but not change system status and settings.
- **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's session can be idle before the user is automatically signed out (from 2 to 720 minutes). The default is 30 minutes.

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" on page 71.

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, start with a hyphen, or include any of the following: " , < \
- **Password**. A password is case sensitive and can have from 8 to 32 characters. If the password contains only printable 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 printable 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 characters. If the password contains only printable 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 printable 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" below
- "Setting the date and time" below
- "Securing the system with FDE" on the next page
- "Setting system properties" on page 59 (cache, disk, scrub, managed logs, and firmware)

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 SMC 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>`.

NOTE You can also access the **Date and Time** panel by selecting on the date and time displayed in the banner.

It is important to set the date and time so that entries in system logs and notifications have correct time stamps. Access the **Date and Time** panel by selecting on the date and time displayed in the banner or by selecting **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 such 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.

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

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, select **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 disk's encrypted data is inaccessible.

ⓘ **IMPORTANT** 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. 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

Select **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.

NOTE 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" on page 38.

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:

- Use **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.
- Use **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.

- "Setting system cache properties" below
- "Setting system disk properties" on the next page
- "Setting system scrub properties" on the next page
- "Setting system managed logs properties" on the next page
- "Setting firmware updates" on page 61

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 tool tips in the panel for specific information regarding each option.

TIP 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 remanufacture

See the tool tips in the panel for specific information regarding each option. For more information about dynamic spares, see "Spares" on page 26.

Setting system scrub properties

The **Scrub Properties** panel (**Settings** > **System** > **Properties** > **Scrub Properties**) enables scrub operations to inspect and fix errors found in disk groups. We recommend that this option be enabled. See the tool tip in the panel for specific information.

TIP If the option is disabled, you can still scrub a selected disk group. See "Disk-group scrub" on page 23.

Creating a schedule for scrub is optional. We recommend that you allow the system to manage the task activity without a schedule.

If **Disk Group Scrub** is enabled and a schedule is not created, the scrubbing process runs at the utility priority level in the background. The scrubbing process is designed to defer to I/O activity in order to avoid impacting system performance. A schedule can be created to further reduce the timeframe when scrubbing is active. The scrubbing task continues to be visible as an in-progress activity on the Dashboard, but the process will only run within the schedule. The scrub process might run outside the scheduled time frame if an inadequate amount of time is allocated to complete the scrub in the system-determined time frame.

When creating a schedule, both a schedule start and schedule stop time are required. Multiple schedules can be created for scrub but care should be taken to avoid overlapping schedules, which could prematurely stop the scrubbing process. A scheduled scrub requires names for the start task and for the stop task as well as the date to begin running the schedule.

The recurrence of the task is set with 3 parameters: frequency, day, and unit:

- Frequency allows for running on the **First**, **Last**, or **Any** day.
- Day allows you to select from: **Day** (meaning every day), **Weekday**, **Weekend Day**, or a specific day of the week.
- Unit allows you to select: **Year** (meaning throughout the year) or a specific month of the year.

After creating a schedule you can modify or delete it.

In order to disable **Disk Group Scrub** you must first delete all schedules for running disk-group scrub tasks.

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 retainment. For more information, see "Managed logs" on page 33. 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 firmware updates

Checkboxes allow for enabling/disabling **Partner Firmware Update** and **Intelligent Disk Firmware Update**. The tooltip next to each checkbox describes its action.

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.

ⓘ **IMPORTANT** We recommend PFU be enabled. Disable this option only if told to do so by a service technician.

When **Intelligent Disk Firmware Update** is enabled (**Settings** > **System** > **Properties** > **Firmware Properties**), disk firmware can be updated using either a manual process or an intelligent disk firmware update process.

NOTE The checkbox only displays on systems enabled with the **Intelligent Disk Firmware Update** option.

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" on page 42.

TIP 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 panel's 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" on page 51.

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 panel's text boxes to receive alert notifications. For details about panel options, see the on-screen tool tips. See "Enabling or disabling system-management services" on page 53 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.

6 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. See the following topics for more information:

- "Storage panel" below
- "Hardware panel" on page 64
- "Firmware panel" on page 65
- "About panel" on page 69
- "Support panel" on page 70

Storage panel

If storage is not yet configured and provisioned for the system, this panel provides configuration options from which you can choose. These options are described in "Configuring and provisioning a new storage system" on page 10.

The **Storage** panel (**Maintenance** > **Storage**) shows the system's storage configuration, including pools, disk groups, and spares (non-ADAPT disk groups only) and enables you to change the configuration. To learn about pools and disk groups, see "Pools" on page 27. To learn about spares, see "Spares" on page 26.

This panel shows a Pool table. In this panel you can:

- View information about disk groups in a pool
- Add a disk group
- Delete a disk group
- Reconfigure storage settings
- Expand an ADAPT disk group
- Scrub a disk group
- View information about disks in a group
- View information about volumes attached to hosts

Viewing information about volumes for each disk group

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 percentage of progress)

In the disk group's 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
- Status
- Target spare capacity (ADAPT)
- Actual spare capacity (ADAPT)
- Stripe width ADAPT (16+3, 16+2, 8+3, and 8+2)
- Interleaved volume count (shown only if interleaved volumes are in use)
- SPEAR

In the disk group's slide-over panel, the **Disks** tab shows information about each disk. Disk location is shown in the format `<enclosure-number>.<disk-slot-number>`. The **Volumes** tab shows information about each volume (name, size, allocated).

Adding a disk group

In the **Maintenance** > **Storage** panel, in the pool where you want to add the disk group, select **Add Disk Group** and follow the on-screen instructions.

TIP To simplify storage system configuration—and make use of key features—choose one of the preconfigured ADAPT protection levels from the wizard, based on your requirements relative to capacity and performance. The (**Maintenance** > **Storage** > **Reconfigure Storage Settings**) action can also be used for this purpose.

When adding a disk group (**Maintenance** > **Storage** > **Pool Configuration** > **Add Disk Group**), you can select a stripe width option from the ADAPT **Stripe Width** dropdown on the **Configuration** panel, provided the disk group has the feature and the protection level is set to ADAPT. The dropdown provides 16+3, 16+2, 8+3, and 8+2 stripe width options for the ADAPT protection level.

Beneath the ADAPT **Stripe Width** dropdown is a **Spare Size** text label with a size field and a units dropdown, which provide for setting the target spare capacity for an ADAPT disk group. Follow the on-screen instructions provided in the tool tip.

For more information about available protection levels, see "RAID levels" on page 18.

IMPORTANT Performance of the system will be reduced until initialization of all disk groups is complete.

Deleting a disk group

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

NOTE Manual configuration is required for deleting individual disk groups. If using **Highest Capacity**, **Highest Performance**, or **Highest Reliability** configurations, use the **Reconfigure Storage Settings** action for this purpose.

Reconfigure storage settings

In the **Maintenance** > **Storage** panel, locate the **Reconfigure Storage Settings** > **delete this configuration and start over** hyperlink if you want to choose another storage configuration. Select the underscored link and follow the on-screen directions.

Expanding an ADAPT disk group

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

ⓘ **IMPORTANT** Validate actual spare capacity according to best practices and set target spare capacity in the CLI prior to expanding an ADAPT disk group. Target spare capacity may only be increased when expanding the disk group.

Scrubbing a disk group

In the **Maintenance** > **Storage** panel, locate the disk group to scrub, display its slide-over panel, select **Scrub Disk Group**, and follow the on-screen directions.

To cancel scrub, select the ⊗ icon.

For more information about the scrub utility, see "Disk-group scrub" on page 23.

Hardware panel

The Hardware panel (**Maintenance** > **Hardware**) shows the system's hardware configuration.

The panel has three sections:

- The top section shows basic information about each enclosure: Rack number, Rack position, Disk slots (used and total).
- The middle section shows a top view or rear view of components located in the enclosure. The dropdown beneath the enclosure view 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's text label updates to show that component, and provides component-specific actions (if available).
- For the enclosure or selected 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.

View	Device	Available actions
Top or Rear	Enclosure	Restart/shutdown system
		Rescan all disks
		Turn on locator LED

Top	Disk (healthy)	Turn on locator LED
	Disk (leftover)	Turn on locator LED
		Clear disk metadata
	Controller module fan (CM fan)	Turn on locator LED
	Disk expander	None
Rear	Power supply	None
	System fan	Turn on locator LED
	Controller module (CM)	Turn on locator LED
	Host port	Reset host port
	Network port	None

This table lists the information shown for a given device.

Device	Information shown
Enclosure	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, GEM A version, GEM B version
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 down count, Temperature, Status, Power on hours, SSD life remaining, FDE state, FDE lock key, FIPS capable, PI formatted, Remanufactured. Disk modules in the main bay are numbered 0 to 95, and disk modules in the controller bay are numbered 96 to 105. Disk modules are accessible from the top of the enclosure.
Power supply	Status, Vendor, Model, Serial number, Revision, Location, Part number, Manufacturing date, Manufacturing location. The two power supply units (PSU), numbered 0 and 1, reside in the PSU slots accessed from the rear of the enclosure.
Controller module (CM)	Controller ID, Locator LED On/Off button, 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 controller modules, labeled as Controller A and B, reside in the controller module slots, and are accessed from the rear of the enclosure.
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)
CM fan module	Module name, Locator LED On/Off button, Location, Module status (OK/not OK), Fan name, Fan status (Up/Down), Fan speed. The two controller fan modules, numbered 4 to 5, reside in the controller bay, and are accessed from the top of the enclosure
Disk expander	Enclosure, Location, Status. The eight SAS expanders, numbered 0 to 7, are labeled as INT EXP in the SMC. They are located in the main bay, and are accessed from the top of the enclosure.
System fan module	Module name, Locator LED On/Off button, Location, Module status (OK/not OK), Fan name (2 fans), Fan Status (Up/Down), Fan speed. The four system fan modules, numbered 0 to 3, reside in the system fan slots accessed from the rear of the enclosure.

For a selected disk module, the **Related Health Actions** dropdown lists additional actions that are accessible depending on the state of the disk. The **Clear Metadata** action is accessible for disks in leftover state. The **Repurpose Secured Disks** and **Import Secured Disks** options are accessiblefor disks that are FDE-capable and in `Secured, Locked` or `Secured, Unlocked` or `Protocol failure` state.

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 may be a newer bundle or an older, previously active bundle.

In this panel you can:

- View information about the system's current firmware bundle
- View whether the system's **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
- Update disk firmware

TIP To aid successful installation and activation of system firmware be sure to read the on-screen directions.

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" on page 68.

Both controllers must run the same firmware version.

In the **System** tab, the process to update firmware is to install firmware bundles and then activate a particular bundle. For a dual-controller system, the following controller firmware-update scenarios are supported:

- Automatic. The partner firmware update (PFU) option is enabled (the default). When you activate controller module firmware on one controller, the firmware is automatically copied over and activated on the partner controller first, and then activated on the current controller.

NOTE We recommend 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 into the partner controller and manually perform the same updates.

Updating controller firmware with the PFU option enabled will ensure that the same firmware version is installed in both controller modules. Access this option by selecting **Settings** > **System** > **Properties** > **Firmware Properties**. PFU uses the following algorithm to determine which controller module will update its partner:

- If both controllers are running the same firmware version, no change is made.
- The controller installed first will send its configuration and settings to the partner controller. Similarly, if a controller is replaced, it will receive configuration information from the partner controller. In both cases, subsequent firmware update behavior for both controllers is determined by the system's unified PFU setting.
- If both controllers were already installed in the system, then the controller with firmware installed first will send its configuration and settings to the partner controller.
- If both controllers are newly installed, then controller A is transferred to controller B.

To install a firmware bundle, follow the on-screen directions and "Best practices for updating firmware" on the next page.

To activate a firmware bundle, select its **Activate this Version** link (**Maintenance** > **Firmware** > **System** > **Firmware Versions** > <firmware-bundle>) to display the **Activate Firmware** dialog and then follow the on-screen directions to enable activation to proceed. As part of the activation process the system will perform 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 SMC login screen will reappear. After you log back in, the **Maintenance** > **Firmware** panel will show that the new firmware is active on the system. An alert will also be generated to inform you that the firmware has been upgraded.

If firmware activation fails, go to **Maintenance** > **Support** > **Collect Logs** and fill in the necessary fields and collect the logs. Such logs will be needed for any support request generated by this failure.

TIP Consider the following points before updating system firmware:

- 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 need to be resolved before attempting to update the firmware. If this problem persists, contact technical support.
- When 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 currently unavailable. When this message is cleared, you may sign in again.
- If PFU is enabled, the amount of time required for updating both controllers is less than 10 minutes.
- If PFU is enabled for the system (**Settings** > **System** > **Properties** > **Firmware Properties** > **Partner Firmware Update** checkbox), 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 and repeat the update. If this problem persists, contact technical support.

Updating disk firmware

IMPORTANT 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.

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 system provides options for manual and intelligent disk firmware updates. To use the manual method, select the **Skip Intelligent Disk Firmware Upgrade** checkbox. To use the intelligent method, clear the checkbox.

NOTE The checkbox only displays on systems enabled with the **Intelligent Disk Firmware Update** option.

Manual disk firmware update

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 on-screen directions and "Best practices for updating firmware" below.

Intelligent disk firmware update

Once the firmware file is uploaded to the Disk Firmware Versions table as Installed/Not Activated—it can be activated using Activate this Version—and in an automated fashion, the system determines which disks can be safely upgraded with the firmware file, and updates them in background. Once completed, the Status column will display Active Version.

To cancel an in-progress intelligent disk firmware update operation, select **Abort disk firmware update**. The batch of disks currently being updated will complete, but no further updates will start. To determine which disks were updated, view the Revision column in the Disk Firmware table.

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" on page 42.

- 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 event 44 in the Event Descriptions Reference Guide and 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. See information about events 172 and 485 in the Event Descriptions Reference Guide.
- 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 selecting **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

Hardware information displays the enclosure number. Expand the table to see the following information for the chassis midplane, the two controller modules, and the eight disk expanders:

- Description
- Part number
- Serial number
- Configuration serial number
- Location

Storage information includes the following for each disk group:

- Disk group name
- Total size
- Available size
- Volumes
- Sector format
- Health

Firmware information includes:

- Bundle version
- 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)

Support panel

The Support panel (**Maintenance** > **Support**) enables you to perform these support-related actions:

- Collect logs
- View the system event history
- View controller module audit logs

For details, see the on-screen directions.

A Other management interfaces

SNMP reference

This topic describes the Simple Network Management Protocol (SNMP) capabilities that Seagate Exos CORVAULT 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, and so on).

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 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.

The system object identifier (`sysObjectID`) for Seagate Exos CORVAULT Storage systems is 1.3.6.1.4.1.347. System uptime is an offset from the first time this object is read.

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 and alerts occurring in the storage system. These events and alerts can be selected by severity and by individual type. A maximum of three SNMP trap destinations can be configured by IP address.

Enterprise event and alert severities include 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/alert ID, the event/alert code type, and a text description generated from the internal event/alert. Equivalent information can also be sent using email or pop-up alerts to users who are logged in to the SMC.

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 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 (FC) host interface ports. Seagate Exos CORVAULT enclosures do not support FC host interface protocol. Hence, the port section of the table is omitted for clarity.

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 SMC. The statistics section is not implemented.

The following table lists the MIB objects, their descriptions and the value set in the storage systems. Unless specified otherwise, objects are *not* settable.

Table 8 FA MIB 2.2 objects, descriptions, and values

Object	Description	Value
`RevisionNumber`	Revision number for this MIB	0220
`UNumber`	Number of connectivity units present	1
`SystemURL`	Top-level URL of the device. For example, `http://10.1.2.3`. If a web server is not present on the device, this string is empty in accordance with the FA MIB2.2 Spec.	Default: `http://10.0.0.1`
`StatusChangeTime`	`sysuptime` timestamp of the last status change event, in centiseconds. `sysuptime` starts at 0 when the Storage Controller boots and keeps track of the up time. `statusChangeTime` is updated each time an event occurs.	0 at startup
`ConfigurationChangeTime`	`sysuptime` timestamp of the last configuration change event, in centiseconds. `sysuptime` starts at 0 when the Storage Controller boots and keeps track of the up time. `configurationChangeTime` is updated each time an event occurs.	0 at startup
`ConnUnitTableChangeTime`	`sysuptime` timestamp of the last update to the `connUnitTable` (an entry was either added or deleted), in centiseconds	0 always (entries are not added to or deleted from the `connUnitTable`)
connUnitTable	**Includes the following objects as specified by the FA MIB2.2 Spec**	

Table 8 FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
`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's 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`
`connUnitConfigurationChangeTime`	Same as `configurationChangeTime`	Same as `configurationChangeTime`
`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
`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 SMC.	Default: info(8)
`connUnitNumEvents`	Number of events currently in the `connUnitEventTable`	Varies as the size of the Event Table varies

Table 8 FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
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 "External details for connUnitRevsTable" on the facing page
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 "External details for connUnitRevsTable" on the facing page
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 connUnitSensorTable" on page 76
connUnitSensorName	Display string containing textual identification of the sensor intended primarily for operator use	See "External details for connUnitSensorTable" on page 76
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 connUnitSensorTable" on page 76
connUnitSensor Characteristic	Characteristics being monitored by this sensor	See "External details for connUnitSensorTable" on page 76
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
connUnitEventIndex	Index into the connectivity unit's event buffer, incremented for each event	Starts at 1 every time there is a table reset or the unit's 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)

Table 8 FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
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
connUnitPortStatFabric Table	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	See the following object entries
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 topic specify values for certain objects described in Table 8.

External details for connUnitRevsTable

Table 9 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 baselevel 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)

Table 9 connUnitRevsTable index and description values (continued)

connUnitRevsIndex	connUnitRevsDescription
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 baselevel 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)
35	Host interface hardware (chip) revision (Controller B)
36	Disk interface hardware (chip) revision (Controller B)

External details for connUnitSensorTable

Table 10 connUnitSensorTable index, name, type, and characteristic values

connUnitSensorIndex	connUnitSensorName	connUnitSensorType	connUnitSensor Characteristic
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)

Table 10 connUnitSensorTable index, name, type, and characteristic values (continued)

connUnitSensorIndex	connUnitSensorName	connUnitSensorType	connUnitSensor Characteristic
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)
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 11 connUnitPortTable index and name values

connUnitPortIndex	connUnitPortName
0	Host Port 0 (Controller A)
1	Host Port 1 (Controller A)
2	Host Port 2 (Controller B)
3	Host Port 3 (Controller B)

Configuring SNMP event notification in the SMC

1. Verify that the storage system's SNMP service is enabled. See "Enabling or disabling system-management services" on page 53.
2. Configure and enable SNMP traps. See "SNMP notifications" on page 61.
3. Optionally, configure a user account to receive SNMP traps. See "Managing SNMPv3 users" on page 56.

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.

In order to view and set system group objects, SNMP must be enabled in the storage system "Enabling or disabling system-management services" on page 53. 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" on page 54. The same users, security protocols, and passwords must be configured in the network management system.

Enterprise trap MIB

To access and download source for this MIB, see www.seagate.com/support/systems/general-support.

Using FTP and SFTP

Although the SMC is the preferred interface for downloading log data and historical disk-performance statistics, you can use SFTP/FTP to do such tasks. SFTP/FTP can also be used for updating firmware and installing security certificates and keys.

NOTE We recommend using SFTP rather than FTP because it is a secured protocol. For this reason, the SFTP/FTP sequence is used when presenting these protocols in tandem.

ⓘ **IMPORTANT** Do not attempt to do more than one of the operations 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
- Each controller's MC logs
- Each controller's event log
- Each controller's debug log
- Each controller's boot log, 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 SFTP/FTP client. A GUI-based SFTP/FTP client might not work.

To download system logs

1. In the SMC, prepare to use SFTP/FTP:
   a. Determine the network-port IP addresses of the system's controllers. See "Configuring controller network ports" on page 51.
   b. Verify that the system's SFTP/FTP service is enabled and take note of the SFTP/FTP service port. See "Enabling or disabling system-management services" on page 53.
   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" on page 54.
2. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the destination directory for the log file.
3. Using the SFTP/FTP 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
   ftp 10.1.0.9
   ```
4. Log in as a user that has permission to use the SFTP/FTP 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 will contain the logs. It is recommended to choose a filename that identifies the system, controller, and date.

   For example:

   ```
   get logs Storage2_A_20120126.zip
   ```
7. In FTP, wait for the message `Operation Complete` to appear. No messages are displayed in SFTP; instead, the `get` command will return once the logs collection is finished.
8. Quit the SFTP/FTP session.

NOTE You must uncompress a zip file before you can view the files it contains. To examine diagnostic data, first view `store_<yyyy>_<mm>_<dd>_<hh>_<mm>_<ss>.logs`.

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 storage system's FTP or SFTP interface 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 a log's data, the log's capacity status is reset to zero to indicate that there is no untransferred data. Log data is controller specific.

For an overview of the log-management feature, see "Managed logs" on page 33.

Use a command-line-based SFTP/FTP client. A GUI-based SFTP/FTP client might not work.

To transfer log data to a log-collection system

1. In the SMC, prepare to use SFTP/FTP:
   a. Determine the network-port IP addresses of the system's controllers. See "Configuring controller network ports" on page 51.
   b. Verify that the system's SFTP/FTP service is enabled. See "Enabling or disabling system-management services" on page 53.
   c. Verify that the user you will log in as has permission to use the SFTP/FTP interface. The same setting allows a user to transfer files using both FTP and SFTP. See "User settings" on page 54.
2. On the log-collection system, 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
   ```

   ```
   ftp 10.1.0.9
   ```
4. Log in as a user with permission to use the SFTP/FTP 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 Storage Controller's four crash logs.
     - `ecdebug`: Expander Controller log.
     - `mc`: Management Controller log.
     - `scdebug`: Storage Controller log.
   - `<filename>` is the file that will contain the transferred data. It is recommended to choose a filename that identifies the system, controller, log type, and date.

   For example:

   ```
   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 will return once the data transfer is finished.

7. Quit the SFTP/FTP session.

NOTE You must uncompress a zip file before you can view the files it contains.

Downloading historical disk-performance statistics

You can access the storage system's SFTP/FTP interface 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 6 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 SFTP/FTP client. A GUI-based SFTP/FTP client might not work.

To retrieve historical disk-performance statistics

1. In the SMC, prepare to use SFTP/FTP:
   a. Determine the network-port IP addresses of the system's controllers. See "Configuring controller network ports" on page 51.
   b. Verify that the system's SFTP/FTP service is enabled. See "Enabling or disabling system-management services" on page 53.
   c. Verify that the user you will log in as has permission to use the SFTP/FTP interface. The same setting allows a user to transfer files using both FTP and SFTP. See "User settings" on page 54.
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`

   `ftp 10.1.0.9`
4. Log in as a user with permission to use the SFTP/FTP interface.
5. Make sure the client is in binary transfer mode. Enter:

   `binary`
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>` is the file that will contain the data. It is recommended to choose a filename that identifies the system, controller, and date.

For example:

```
get perf:start.2012-01-26.12:00.PM.end.2012-01-26.23:00.PM Storage2_A_20120126.csv
```

7. 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.
8. Quit the SFTP/FTP session.

Updating firmware

As a user in the `manage` role, you can update the versions of firmware in controller modules and disks.

TIP 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.

IMPORTANT Consider the following points before performing a firmware update:

- If a disk group is quarantined, resolve the problem that is causing the disk group to be quarantined before updating firmware. See information about events 172 and 485 in the Event Descriptions 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 event 44 in the Event Descriptions Reference Guide and information about the `clear cache` command in the CLI Reference Guide.
- If the system's health is `Fault`, firmware update will not proceed. Before you can update firmware, you must resolve the problem specified by the Health Reason value in the System Overview panel.

Updating controller-module firmware

In a dual-controller system, both controllers should run the same firmware version. 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 partner controller's IP address 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 product's Release Notes.

To update controller module firmware

1. Obtain the appropriate firmware file and download it to your computer or network.
2. In the SMC, prepare to use SFTP/FTP:

   a. Determine the network-port IP addresses of the system's controllers.
   b. Verify that the system's SFTP/FTP service is enabled.
   c. Verify that the user you will log in as has permission to use the SFTP/FTP interface. The same setting allows a user to transfer files using both FTP and SFTP.

3. If the storage 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 navigate 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
   ftp 10.1.0.9
   ```
6. Log in as an SFTP/FTP user.
7. Make sure the client is in binary transfer mode. Enter:

   ```
   binary
   ```
8. 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 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 SFTP/FTP 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.

   NOTE If you are using a Windows SFTP/FTP client, during firmware update a client-side SFTP/FTP application issue or time out setting can cause the SFTP/FTP session to be aborted. If this issue persists try using the SMC to perform the update, use another client, or use another SFTP/FTP application.

   If the Storage Controller cannot be updated, the update operation is canceled. If the SFTP/FTP prompt does not return, quit the SFTP/FTP 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 SFTP/FTP session to the local MC is closed. You will need to monitor the system using a management interface to determine when the update is complete. If PFU is enabled, both controllers will update. If the SMC 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.
9. Quit the SFTP/FTP session.

10. Clear your web browser's cache, then sign in to the SMC. 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. After firmware update has completed on both controllers, 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 and repeat the update. If this problem persists, contact technical support.

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.

To prepare for update

1. Obtain the appropriate firmware file and download it to your computer or network.
2. Check the disk manufacturer's documentation 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, continue with the next step. Otherwise, in the SMC, 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 SMC, prepare to use SFTP/FTP:
    a. Determine the network-port IP addresses of the system's controllers.
    b. Verify that the system's SFTP/FTP service is enabled.
    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.
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.

To update disk firmware

1. As a user with the `manage` role, open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the directory containing the firmware file to load.
2. Enter:

    `sftp -P` `<port> <controller-network-address>` or **`ftp`** `<controller-network-address>`

    For example:

    ```
    sftp -P 1022 10.235.216.152
    ftp 10.1.0.9
    ```
3. Log in as an SFTP/FTP user.

4. Make sure the client is in binary transfer mode. Enter:

   binary

5. Either:

   - To update all disks of the type that the firmware applies to, enter:

     put <firmware-file> **disk**

   - To update specific disks, enter:

     put <firmware-file> **disk:**<enclosure-ID>**:**<slot-number>

   For example:

   put <firmware-file> **disk:1:11**

   ⚠ **CAUTION** Do not power cycle an enclosure 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 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.

6. If you are updating specific disks, repeat the previous step for each remaining disk to update.
7. Quit the SFTP/FTP session.
8. If the updated disks must be power cycled:

   a. Shut down both controllers by using the SMC.

   b. Power cycle the storage system as described in your product's Hardware Installation and Maintenance Guide.
9. 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.

You can use SFTP/FTP to install a device certificate signed by a Certificate Authority (CA). The process is:

1. Use the CLI to create the Certificate Signing Request (CSR), copy the output of the command (an encoded CSR) to a file, and sign the CSR using a CA.
2. Use SFTP/FTP to upload each certificate of the CA trust chain and then to upload the signed certificate. The CA certificate (s) must be uploaded first so they can be used to verify the signed device certificate.

NOTE The SFTP/FTP interface supports uploading only one trust and one user certificate. To upload multiple certificates, use the `add certificate` API command as described in the CLI Reference Guide.

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.

To install a security certificate

1. In the SMC, prepare to use SFTP/FTP:
   a. Determine the network-port IP addresses of the system's controllers. See "Configuring controller network ports" on page 51.
   b. Verify that the system's SFTP/FTP service is enabled. See "Enabling or disabling system-management services" on page 53.
   c. Verify that the user you will log in as has permission to use the SFTP/FTP interface. The same setting allows a user to transfer files using both FTP and SFTP. See "Managing local users" on page 55.
2. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate 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

   ftp 10.1.0.9
4. Log in as a user that has permission to use the SFTP/FTP interface.
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>: Name of the certificate file for your specific system.
   - usr: Device/server certificate uploaded by end user.
   - mfg: Reserved for use in engineering mode only.

Using SLP

Seagate Exos CORVAULT 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.

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 shown in Table 12 and populate the configuration attributes shown in Table 13.

Table 12 Interfaces advertised by SLP

Interface (protocol) description	Advertisement string
HTTP	`service:api:http`
HTTPS	`service:api:https`
Telnet	`service:ui:telnet`
SSH	`service:ui:ssh`
SFTP/FTP (firmware upgrade)	`service:firmware-update:sftp/ftp`
SNMP	`service:api:snmp`

Table 13 SLP attributes shown for a storage system

SLP attribute	Corresponding property shown by the CLI `show system detail` command in XML API mode
`x-system-name`	`system-name`
`x-system-contact`	`system-contact`
`x-system-location`	`system-location`
`x-system-information`	`system-information`
`x-midplane-serial-number`	`midplane-serial-number`
`x-vendor-name`	`vendor-name`
`x-product-id`	`product-id`
`x-product-brand`	`product-brand`
`x-wwnn`	`current-node-wwn`
`x-platform-type`	`platform-type`
`x-bundle-version`	no corresponding property
`x-build-date`	no corresponding property
`x-mac-address`	no corresponding property
`x-top-level-assembly-part-number`	no corresponding property
`x-top-level-assembly-serial-number`	no corresponding property

You can enable or disable the SLP service in the SMC, as described in "Enabling or disabling system-management services" on page 53, 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 openSLP.org.

Using SeaCombine

Seagate Exos CORVAULT Storage systems support SeaCombine: a Seagate specific command line software utility for retrieving logs from enclosures with HTTPS and SFTP functionality. After pulling and parsing a specific log, SeaCombine saves a binary and JSON file for a specific enclosure drive slot.

NOTE For information about using SeaCombine for downloading drive log data, see the *Seagate SeaCombine User Guide*.

B 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" on page 33.

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.

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*:

- In 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 will be uniquely named to indicate the log-file type, the date/time of creation, and the storage system. This information will also be in the email subject line. The file name format is `logtype_<yyyy>_<mm>_<dd>_<hh>_<mm>_<ss>.zip`.
- In pull mode, when log data has accumulated to a significant size, the system sends notification events via email or SNMP traps, to the log-collection 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. The storage system's SFTP/FTP interface can be used to transfer the appropriate logs to the log-collection system, as described in "Transferring log data to a log-collection system" on page 80.

Log file details

- 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:

```
\Storage1
  \crash
  \ecdebug
  \mc
  \scdebug
\Storage2
  \crash
  \ecdebug
  \mc
  \scdebug
```

In push mode, when the administrator receives an email with an attached EC debug 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 storage system's SFTP/FTP interface to get the log and save it into the `scdebug` subdirectory of the Storage2 directory.

C Settings changed by restoring defaults

This page summarizes the system settings that result from using the CLI `restore defaults` command.

Table 14 System information settings

Setting	Value
System name	Uninitialized Name
System contact	Uninitialized Contact
System location	Uninitialized Location
System information	Uninitialized Info

Table 15 Management protocols settings

Setting	Value
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

Table 16 User settings

Setting	Value
Users	Users are preserved.

Table 17 Miscellaneous settings (1)

Setting	Value
CLI session timeout	Preserved
Tasks and schedules	Preserved
Management Controller debug logs	Preserved
Management Controller event logs	Preserved
Storage Controller debug logs	Preserved
Storage Controller event logs	Preserved
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

Table 18 SNMP settings

Setting	Value
SNMP trap notification level	None
SNMP trap host IPs	0.0.0.0
SNMP read community	`public`
SNMP write community	`private`

Table 19 SMTP settings

Setting	Value
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
Event notification	None
Alert notification	Enabled
Proxy setting	Cleared

Table 20 LDAP settings

Setting	Value
LDAP parameters	Cleared
LDAP settings	Disabled (server IP defaulted to 0.0.0.0)
User groups	Preserved
Audit log	Preserved

Table 21 Syslog settings

Setting	Value
Syslog parameters	Cleared
Syslog settings	Disabled (host IP defaulted to 0.0.0.0)

Table 22 Miscellaneous settings (2)

Setting	Value
SSL/SSH certificates	Preserved
Disk group metadata	Preserved

Table 23 Advanced settings

Setting	Value
Disk group background scrub	Enabled
Partner firmware upgrade	Enabled

Table 23 Advanced settings (continued)

Setting	Value
Utility priority	High
SMART	Enabled
Dynamic spare configuration	Disabled
Enclosure polling rate	5 seconds
Host control of caching	Disabled
Sync cache mode	Immediate
Missing LUN response	Illegal Request
Controller failure	Disabled
Supercap failure	Enabled
Power supply failure	Disabled
Fan failure	Disabled
Temperature exceeded	Disabled
Partner notify	Disabled
Auto write back	Enabled
Managed logs	Disabled
Auto stall recovery	Enabled (for failover/failback, not I/O)
Restart on CAPI fail	Enabled
Slot affinity	Disabled
Remanufacture	Enabled
Disk protection information	Disabled
Slow disk detection	Enabled
Scrub schedule	Disabled

Table 24 FDE settings

Setting	Value
FDE settings	Preserved

Table 25 Enclosure settings

Setting	Value
Name	Cleared
Location	Cleared
Rack number	0
Rack position	0

Table 26 Host settings

Setting	Value
Host and initiator nicknames and profiles	Preserved
Host groups	Preserved

Table 27 Volume settings

Setting	Value
Volume identifying information	Preserved

Table 28 Miscellaneous settings (3)

Setting	Value
CLI parameters	Configured users remain unchanged.
Debug log parameters	Each parameter is reset to its default as documented for the `set debug-log-parameters` CLI command.
Volume cache settings	Preserved
Expander PHY settings	Controller module root expander PHY settings are cleared
Hedged reads	Enabled (wait time up to 1.5s)

D System configuration limits

Table 29 System configuration limits

Feature	Value
Enclosures and disks	
Maximum enclosures per system	1
Maximum disks per system	106
Maximum disk size	26TB
Disk groups	
Storage model	Linear
Maximum non-ADAPT disk-group size	Limited by the number and capacity of disks supported by the RAID level applied to the disk group
Maximum disk groups per controller module	32
Minimum/maximum disks per disk group	RAID 1: 2/2 RAID 5: 3/16 RAID 6: 4/16 RAID 10: 4/16 ADAPT: 12/106
Adjustable spare capacity	Up to 30% of available capacity
Maximum dedicated spares per disk group	4
Maximum global spares per system	64
Maximum ADAPT disk groups per controller module	4
Maximum ADAPT disk group size	1.5PiB
ADAPT stripe width (data+parity)	8+2, 8+3, 16+2, 16+3
Volumes, initiators, hosts, and mapping	
Maximum volumes per system	1024 (512 recommended)
Maximum volume (LUN) size	1.5PiB
Maximum mappable volumes (LUNs) per disk group	128
Maximum mappable volumes (LUNs) per controller module	128
Maximum mappable volumes (LUNs) per controller module	512
Maximum volumes per controller module	1024
Maximum volumes per host port	1024 (Microsoft Windows limits access to 256)
Maximum initiators per host port	1024
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

Table 29 System configuration limits (continued)

Feature	Value
Maximum commands per LUN (preferred path)	1024 per port
Maximum queue depth per host port	1024
Maximum SAS host-port link speed	12Gb
Miscellaneous	
Maximum SCSI reservations per system	1024
Maximum SCSI reservations per LUN	1
Maximum SCSI registrations per system	32768
Maximum SCSI registrations per LUN	SAS: 85

E Scripting reference for CLI setup commands

This appendix describes underlying CLI commands used by the SMC for configuring storage. CLI commands used to set maximum capacity, maximum reliability, and maximum performance configurations are diagrammed with examples to help facilitate scripting and configuring of multiple storage systems in parallel. Command examples and their components are described in **Highest Capacity**, **Highest Reliability**, and **Highest Performance** sections. These configurations apply to enclosures equipped with SAS or fixed-capacity SATA drives.

TIP Before using CLI commands for storage system setup, become familiar with the CLI Reference Guide:

- Read "Using the CLI > Command syntax" (entire subsection: study specifying disks; disk groups; and volumes).
- Read "Using a script to access the CLI" (entire subsection).

Highest Capacity

```
add disk-group type linear level adapt disks 0.0,0.1,0.2,0.3,0.4,0.5,0.6,0.7,0.8,0.9,0.10,0.11,0.24,
0.25,0.26,0.27,0.28,0.29,0.30,0.31,0.32,0.33,0.34,0.35,0.48,0.49,0.50,0.51,0.52,0.53,0.54,0.55,0.56,
0.57,0.58,0.59,0.72,0.73,0.74,0.75,0.76,0.77,0.78,0.79,0.80,0.81,0.82,0.83,0.96,0.97,0.98,0.99,0.100
stripe-width 16+2 spare-capacity 16.0TB mode online interleaved-volume-count 8 interleaved-basename
Volume_ "dg01"
```

```
add disk-group type linear level adapt disks 0.12,0.13,0.14,0.15,0.16,0.17,0.18,0.19,0.20,0.21,0.22,
0.23,0.36,0.37,0.38,0.39,0.40,0.41,0.42,0.43,0.44,0.45,0.46,0.47,0.60,0.61,0.62,0.63,0.64,0.65,0.66,
0.67,0.68,0.69,0.70,0.71,0.84,0.85,0.86,0.87,0.88,0.89,0.90,0.91,0.92,0.93,0.94,0.95,0.101,0.102,
0.103,0.104,0.105 stripe-width 16+2 spare-capacity 16.0TB mode online interleaved-volume-count 8
interleaved-basename Volume_ "dg02"
```

Command item	Description of item: command/ command parameter/ parameter value
add disk-group	CLI command given in the example: Add a disk-group using specified disks. See also CLI Reference Guide or on-line help for the full command description.
type	Storage type parameter: linear
level	Storage protection level: ADAPT
disks	Disk IDs to include in disk-group: enclosure number followed by disk slot. Use a comma-separated list of disk IDs. Splitting them up rather than just using the first 53 for one group and the last 53 for the other is deliberate for the command example shown.
stripe-width	16+2: Each stripe contains 16 data chunks and 2 parity chunks.
spare-capacity	16.0TB: Default is 2 disks, so the size is 16TB because this system has 8TB drives. On an 18TB drive system this would default to 36.0TB.
mode	online: Bring the disk-group online rather than setting it to offline.
interleaved-volume-count	8: Create 8 optimally sized volumes on the disk-group.
interleaved-basename	Volume_: Prefix the auto-numbered volume names with Volume_ , e.g., Volume_0000, Volume_0001, etc.
"dg01"|"dg02"	The names of the disk-groups used in the examples

NOTE There is no separate volume creation command for **Highest Capacity** because the volumes are auto-created with the correct size by specifying the interleaved-volume-count.

Highest Performance

Disk Groups

```
add disk-group type linear level adapt disk 0.0,0.1,0.2,0.3,0.4,0.5,0.6,0.7,0.8,0.9,0.10,0.11,0.12
assigned-to a stripe-width 8+2 spare-capacity 8.0TB mode online "dg01"
```

```
add disk-group type linear level adapt disks 0.13,0.14,0.15,0.16,0.17,0.18,0.19,0.20,0.21,0.22,0.23,
0.24,0.25 assigned-to a stripe-width 8+2 spare-capacity 8.0TB mode online "dg02"
```

```
add disk-group type linear level adapt disks 0.26,0.27,0.28,0.29,0.30,0.31,0.32,0.33,0.34,0.35,0.36,
0.37,0.38 assigned-to a stripe-width 8+2 spare-capacity 8.0TB mode online "dg03"
```

```
add disk-group type linear level adapt disks 0.39,0.40,0.41,0.42,0.43,0.44,0.45,0.46,0.47,0.48,0.49,
0.50,0.51,0.52 assigned-to a stripe-width 8+2 spare-capacity 8.0TB mode online "dg04"
```

```
add disk-group type linear level adapt disks 0.53,0.54,0.55,0.56,0.57,0.58,0.59,0.60,0.61,0.62,0.63,
0.64,0.65 assigned-to b stripe-width 8+2 spare-capacity 8.0TB mode online "dg05"
```

```
add disk-group type linear level adapt disks 0.66,0.67,0.68,0.69,0.70,0.71,0.72,0.73,0.74,0.75,0.76,
0.77,0.78 assigned-to b stripe-width 8+2 spare-capacity 8.0TB mode online "dg06"
```

```
add disk-group type linear level adapt disks 0.79,0.80,0.81,0.82,0.83,0.84,0.85,0.86,0.87,0.88,0.89,
0.90,0.91 assigned-to b stripe-width 8+2 spare-capacity 8.0TB mode online "dg07"
```

```
add disk-group type linear level adapt disks 0.92,0.93,0.94,0.95,0.96,0.97,0.98,0.99,0.100,0.101,
0.102,0.103,0.104,0.105 assigned-to b stripe-width 8+2 spare-capacity 8.0TB mode online "dg08"
```

Command item	**Description of item: command/ command parameter/ parameter value**
add disk-group	CLI command given in the example: Add a disk-group using specified disks. See also CLI Reference Guide or on-line help for the full command description.
type	Storage type parameter: linear
level	Storage protection level: ADAPT
disks	Disk IDs to include in disk-group: enclosure number followed by disk slot. Use a comma-separated list of disk IDs.
assigned-to	a|b: The controller module to which the disk-groups is assigned. The first 4 disk-groups are assigned to A; and the last 4 disk-groups are assigned to B.
stripe-width	8+2: Each stripe contains 8 data chunks and 2 parity chunks.
spare-capacity	8.0TB: Default is 1 disk, so the size is 8TB because this system has 8TB drives. On an 18TB drive system this would default to 18.0TB.
mode	online: Bring the disk group online rather than setting it to offline.
"dg01"–"dg08"	The names of the disk-groups used in the examples

TIP Before using CLI commands for storage system setup, read this content:

- For information about creating `disk-groups`, see "Linear disk group configuration" on page 17.
- For information about RAID levels, see "RAID levels" on page 18 and "ADAPT" on page 20.
- Within the CLI Reference Guide, see the `add disk-group` command topic.

Volumes

Command examples and their components are described. For **Highest Performance**, the system issues separate commands to create the `volumes` because it uses non-interleaved `disk-groups`. The system creates all the `disk-groups` first, then it creates the `volumes`:

```
create volume pool "00c0ff4720600000d7862e6000000000" size 149854093312 "Volume_0001"
```

There are 8 `create volume` commands, but they.are all the same apart from the `size` parameter for the 4th and 8th disk groups, as they have 14 disks rather than 13.

Command item	Description of item: command/ command parameter/ parameter value
`create volume`	CLI command given in the example: Create a `volume` in a `pool`. See also CLI Reference Guide or on-line help for the full command description.
`pool`	"00c0ff4720600000d7862e6000000000": The system uses serial number; but use of the `disk-group` name is allowed (e.g., dg01).
`size`	149854093312: Maximum `volume size` in blocks for the `disk-group` Interactive method: • Log in to the CLI • Run `set cli-parameters json` (shows output in JSON format) • Run `show disk-groups` (note the `size-numeric` property output for the `disk-group`) Script method: Get the `size-numeric` property from the `/api/show/disk-groups/` end point (see CLI Reference Guide > "Using a script to access the CLI")
"Volume_0001"	The name of the `volume` used in the example

TIP Before using CLI commands for storage system setup, read this content:

- For information about `pools` and `volumes`, see "Pools" on page 27 and "Volumes" on page 27.
- For information about provisioning volumes, see "Working with volumes" on page 45.
- Within the CLI Reference Guide, see the `create volume` command topic.

Highest Reliability

Disk Groups

```
add disk-group type linear level adapt disks 0.0,0.1,0.2,0.3,0.4,0.5,0.6,0.7,0.8,0.9,0.10,0.11,0.24,
0.25,0.26,0.27,0.28,0.29,0.30,0.31,0.32,0.33,0.34,0.35,0.48,0.49,0.50,0.51,0.52,0.53,0.54,0.55,0.56,
0.57,0.58,0.59,0.72,0.73,0.74,0.75,0.76,0.77,0.78,0.79,0.80,0.81,0.82,0.83,0.96,0.97,0.98,0.99,0.100
stripe-width 16+3 spare-capacity 24.0TB mode online "dg01"
```

```
add disk-group type linear level adapt disks 0.12,0.13,0.14,0.15,0.16,0.17,0.18,0.19,0.20,0.21,0.22,
0.23,0.36,0.37,0.38,0.39,0.40,0.41,0.42,0.43,0.44,0.45,0.46,0.47,0.60,0.61,0.62,0.63,0.64,0.65,0.66,
0.67,0.68,0.69,0.70,0.71,0.84,0.85,0.86,0.87,0.88,0.89,0.90,0.91,0.92,0.93,0.94,0.95,0.101,0.102,
0.103,0.104,0.105 stripe-width 16+3 spare-capacity 24.0TB mode online "dg02"
```

Command item	Description of item: command/ command parameter/ parameter value
add disk-group	CLI command given in the example: Add a disk-group using specified disks. See also CLI Reference Guide or on-line help for the full command description.
type	Storage type parameter: linear
level	Storage protection level: ADAPT
disks	Disk IDs to include in disk-group: enclosure number followed by disk slot. Use a comma-separated list of disk IDs. Splitting them up rather than just using the first 53 for one group and the last 53 for the other is deliberate for the command example shown.
stripe-width	16+3: Each stripe contains 16 data chunks and 3 parity chunks.
spare-capacity	24.0TB: Set the target spare capacity to the sum of the size of two disks in the disk group, which is sufficient to fully recover fault tolerance after loss of any two disks in the group.
mode	online: Bring the disk-group online rather than setting it to offline.
"dg01"|"dg02"	The names of the disk-groups used in the examples.

TIP Before using CLI commands for storage system setup, read this content:

- For information about creating disk-groups, see "Linear disk group configuration" on page 17.
- For information about RAID levels, see "RAID levels" on page 18 and "ADAPT" on page 20.
- Within the CLI Reference Guide, see the add disk-group command topic.

Volumes

Command examples and their components are described. For **Highest Reliability**, the system issues separate commands to create the volumes because it uses non-interleaved disk-groups. The system creates all the disk-groups first, then it creates the volumes:

```
create volume pool A size 14GB "Volume_0001"
```

There are 2 create volume commands, but they.are both the same.

Command item	Description of item: command/ command parameter/ parameter value
create volume	CLI command given in the example: Create a volume in a pool. See also CLI Reference Guide or on-line help for the full command description.
pool	Use the pool name (A). You can also use the serial number or disk-group name for linear storage (e.g., dg01).
size	14GB: Maximum volume size for the disk-group Interactive method: • Log in to the CLI • Run set cli-parameters json (shows output in JSON format) • Run show disk-groups (note the size-numeric property output for the disk-group) Script method: Get the size-numeric property from the /api/show/disk-groups/ end point (see CLI Reference Guide > "Using a script to access the CLI")
"Volume_0001"	The name of the volume used in the example

TIP Before using CLI commands for storage system setup, read this content:

- For information about `pools` and `volumes`, see "Pools" on page 27 and "Volumes" on page 27.
- For information about provisioning volumes, see "Working with volumes" on page 45.
- Within the CLI Reference Guide, see the `create volume` command topic.

F Multipath configuration

This appendix describes multipath configuration for the Exos CORVAULT storage system.

Seagate 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.

To enable 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**, select **Role-based** or **Feature Based Installation**.
4. Select **Next**.
5. Select the server from the pool and then select **Next**.
6. Select **Next** again to go to the feature selection window.
7. Select the **Multipath IO** check box and then select **Next**.
8. Select **Install**.
9. When installation is complete, select **Close**.
10. Reboot the storage system using one of the restart methods provided (automatic or manual).
11. In the Server Manager Tools menu, select **MPIO**.
12. Select the **Discover Multi-Paths** tab.
13. Select check boxes for the devices you want to support, and then select **Add**.
14. When prompted, reboot the system.

    When the reboot is complete, MPIO is ready to use.

To enable MPIO on Windows Server 2019 (Standard)

1. Start Server Manager if it is not already running.
2. From the **Server Manager Dashboard** > **Add Roles and Features** wizard, select **Server Roles**.
3. In the **Roles** list select **File and Storage Services** > **Storage Services**.
4. Select **Tools** > **MPIO**. Follow the dialog's tabs to complete MPIO enablement.

Alternatively, once iSCSI is enabled from the Server Manager, you can select the iSCSI initiator icon from the taskbar, and follow the steps outlined in the iSCSI Initiator Properties dialog.

1. Select the **Targets** tab and confirm Name and Status (Connected).
    a. Within the Discovered targets pane, select the target.
    b. Select the **Devices** button.
2. Within **Devices** > **Multipath IO**, select the **MPIO** button.
3. Within the MPIO tab:

a. Select load balance policy from the dropdown.

b. View the policy description.

c. View or edit device paths.

To enable MPIO on Linux

This procedure is general in nature, and does not address specific nuances of different Linux operating systems.

1. Ensure that the multipath daemon is installed and set to start at run-time. Linux command:

```
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 Seagate system. The following table specifies attributes that should be set.

NOTE Attributes shown in the table example pertain to RHEL 8.6 and may change with each major version.

Attribute	Value
vendor	<vendor-name>
product	<product-ID>
path_grouping_policy group	group_by_prio
uid_attribute	"ID_SERIAL"
prio	alua
path_selector	"queue-length 0"
path_checker	tur
failback	immediate
no_path_retry	5
alias_prefix	"mpath"

To discover multipath capable devices and obtain the vendor and product ID values, run the following commands.

Discover multipath capable devices: `[root@smc10 device]# lsblk`

Example (`sdag`/`sdah`) output from the command is shown in the following table.

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPNT
sdag	66:0	0	83.4T	0	disk	
3600c0ff000f48b2944a33c6307000000	253:0	0	83.4T	0	mpath	
sdah	66:16	0	83.4T	0	disk	
3600c0ff000f48b2944a33c6308000000	253:2	0	83.4T	0	mpath	

Get vendor name: `[root@smc10 device]# cat /sys/block/sdag/device/vendor`

```
SEAGATE
```

Get product ID: `[root@smc10 device]# cat /sys/block/sdag/device/model`

```
6575
```

NOTE Within the blacklist section of the `/etc/multipath.conf` file, Seagate strongly recommends using WWID values, found in `/etc/multipath/bindings`, not WWN values.

3. Instruct the multipath daemon to reload the `multipath.conf` file. Linux command:

   ```
   systemctl restart multipathd
   ```

4. Determine if the multipath daemon used ALUA to obtain the optimal/non-optimal paths. Linux command:

   ```
   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

4U106

An enclosure that is four rack units in height and can contain 106 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.

ADR

Autonomous Drive Regeneration (ADR). HDD technology that detects and removes a bad head and its corresponding media surface, and returns the disk to service minus the capacity of the bad surface. ADR is also known as remanufacture.

AES

Advanced Encryption Standard.

array

Synonym for 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.

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 (AWT). A setting that specifies when the RAID controller cache mode automatically changes from write-back to write-through.

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.

chassis

The sheet metal housing of an enclosure.

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.

CMA

Cable management arm (CMA).

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 can contain 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 supercapacitor pack and nonvolatile 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.

dedicated spare

A disk that is reserved for use by a specific linear disk group to replace a failed disk.

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 that provides storage capacity. The number of disks that a disk group can contain is determined by its RAID level.

DNS

Domain Name System.

DSP

Digital signal processor.

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.

E2E PI

End-to-end protection information (E2E PI). A feature that consists of two discrete processes: host PI and disk PI. Host PI primarily verifies the integrity of data sent from the host server during transport by checking protection field information attached to the data. It also adds protection fields to data requested by the host server, so the host server can perform the process in reverse. Disk PI adds new protection fields to data sent from the host server so that it can check the integrity of the stored data when retrieved. Host and disk PI are separately enabled and operate independently. You can enable one, both, or neither of these processes. See also PI.

EC

Expander Controller (EC). A processor (located in the SAS expander in each IOM) that controls the SAS expander and provides SES functionality. See also EMP.

EEPROM

Electrically erasable programmable ROM.

eMMC

Electro-magnetic memory card. Also referred to as memory card, non-volatile memory.

EMP

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.

enclosure

A physical storage device that contains I/O modules, disks, 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.

End-to-End Protection Information

End-to-end protection information (E2E PI). A feature that consists of two discrete processes: host PI and disk PI. Host PI primarily verifies the integrity of data sent from the host server during transport by checking protection field information attached to the data. It also adds protection fields to data requested by the host server, so the host server can perform the process in reverse. Disk PI adds new protection fields to data sent from the host server so that it can check the integrity of the stored data when retrieved. Host and disk PI are separately enabled and operate independently. You can enable one, both, or neither of these processes. See also PI.

ESD

Electrostatic discharge.

Expander Controller

Expander Controller (EC). A processor (located in the SAS expander in each IOM) that controls the SAS expander and provides SES functionality. See also EMP.

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.

FDE

Full Disk Encryption (FDE). 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.

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.

hedged read

A method to reduce latency for read requests by using RAID parity information to reconstruct requested data.

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 set of hosts for ease of management, such as for volume-attachment operations.

host port

A port on a controller module that interfaces to a server, 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 is a controller 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.

interleaved volume

A method of disk storage with ADAPT disk groups that puts sequential information into nonsequential disk sectors, which results in faster read times.

IOM

Input/output module (I/O module, IOM). An IOM is a controller module.

IOPS

I/O operations per second.

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.

JSON

JavaScript Object Notation providing a human readable representation of data that is also easily handled by computer applications.

LBA

Logical block address. The address used for specifying the location of a block of data.

LDAP

Local directory access protocol.

LDAPS

LDAP over SSL.

leftover

The state of a disk that the system has excluded from a disk group because the timestamp in the disk's metadata 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 disk's metadata 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.

lock key

A system-generated value that manages the encryption and decryption of data on FDE-capable disks. See also FDE, passphrase.

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.

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 (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.

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 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. Synonyms for this action include present and map.

network port

The Ethernet port on a controller module through which its Management Controller is connected to the network.

NRAID

Non-RAID, nonstriped mapping to a single disk.

NTP

Network time protocol.

OData

The Open Data protocol is a set of common RESTful conventions providing interoperability between APIs.

OID

Object Identifier. In SNMP, an identifier for an object in a MIB.

onboarding

A wizard-based process in the Storage Management Console that, after preboarding is complete, leads you through configuring and provisioning a new system.

orphan data

See unwritable cache data.

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.

PDU

Power distribution unit. The rack power-distribution source to which a PSU connects.

PFU

Partner firmware update (PFU). 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.

PI

Protection Information. An 8-byte protection field for each 512-byte sector of data that consists of a Guard, Reference, and Application tag. See also E2E PI.

POST

Power-on self test. Tests that run immediately after a device is powered on.

preboarding

A wizard-based process in the Storage Management Console that leads you through preparing a new system for use. Preboarding is followed by onboarding.

provisioning

The process of creating storage volumes, mapping them to initiators or hosts, and configuring data-protection options.

PSU

Power supply unit FRU.

RAID head

Synonym for controller enclosure.

RBOD

"RAID bunch of disks." See controller enclosure.

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.

Redfish

RESTful API and Data model to manage network devices and services maintained by the DMTF organization.

remanufacture

Autonomous Drive Regeneration (ADR). HDD technology that detects and removes a bad head and its corresponding media surface, and returns the disk to service minus the capacity of the bad surface. ADR is also known as remanufacture.

repurpose

A method by which all data in a storage 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.

RESTful

An architectural pattern to API design that provides inherent use of HTTP primitives operating on a data model.

SAS

Serial Attached SCSI (Small Computer System Interface).

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.

SED

Self-encrypting drive. A disk drive that provides hardware-based data encryption and supports use of the storage system's 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.

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 tasks such as 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.

shipkit

The package containing an enclosure, accessories, and related materials that a customer receives.

SLAAC

Stateless address autoconfiguration.

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.

SMC

Storage Management Console 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.

SPEAR

Stripe Parity Error Assessment on Read. A data-protection feature for disk groups that ensures data integrity by verifying consistency of the stripe with each read access.

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.

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 Management Console

Storage Management Console 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.

storage system

A controller enclosure. Product documentation and interfaces use the terms storage system and system interchangeably.

Swordfish

Extends the Redfish API to manage storage equipment and storage services. Maintained by the SNIA organization.

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.

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.

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's 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.

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 SMC and CLI interfaces.

volume

A logical representation of a fixed-size, contiguous span of storage that is presented to host systems for the purpose of storing data.

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

Storage Management Console 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.

Index