HPE MSA 2070/2072 Storage Management Guide

Abstract

This guide covers concepts used to administer an HPE MSA 2070/2072 Storage system and provides information about managing the system by using its web interface, the Storage Management Utility (SMU).

Revision history

20-MSAG7-SMG-ED1	December 2024
Initial HPE release.	

Contents

Figures

Tables

1 Getting started

The Storage Management Utility (SMU) is a web-based application for configuring, monitoring, and managing the storage system. The SMU 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 SMU. 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 SMU, 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 *HPE MSA 2070/2072 CLI Reference Guide*.

Product information that is specific to a product model is identified by model name in the text. Otherwise, the content of this guide applies to all models.

Product features

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

- **Storage system setup**: the capability to initially connect to a system using the SMU, 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.
- **MSA-DP+ data protection**: RAID-based data protection level that emphasizes efficiency, as described in "MSA-DP+" on page 23.
- **Replication**: the capability to replicate block-level data from a volume in a primary system to a volume in a secondary system, as described in "Replication" on page 46.
- **Update firmware**: the capability to notify users of available firmware updates to controller modules, expansion modules, and disk modules with newer/compatible firmware versions as they become available, as described in "Updating firmware" on page 39.
- **Performance metrics**: the ability to monitor storage system performance and statistics via data-driven graphing—displaying dynamic or historical metrics—as described in "Performance panel" on page 61.
- **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 59.
- **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 41.
- **SSDs**: the ability to use solid-state drives to enhance storage system performance, as described in "SSDs" on page 25. The ability to move data automatically from one class of disks to another—based on data access patterns—is described in "Automated tiered storage" on page 32.
- **Virtual storage**: a storage model that maps logical components to physical media—using paged-storage technology—to virtualize data storage, as described in "Virtual storage" on page 19.
- **IPv6 support**: the capability for the storage system to support IPv6 (Internet Protocol version 6) functionality—either exclusively or in addition to IPv4—as described in "Configuring controller network ports" on page 79.
- **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.
- 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).
- **MC Auxiliary Services (MAS)**: a feature that provides microservices-based architectural support. The services, called MAS agents or MAS applications, can be installed using a signed encrypted agent bundle for independent feature delivery of future functionalities like diagnostics, telemetry and more. This feature employs a MAS Manager running in the Management Controller to manage and interface with the MAS service agents using an OpenAPI specification.

Configuring and provisioning a new storage system

When you connect to the system for the first time, a wizard in the SMU 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, password, and select a language (once complete, you will be logged into the system as this user)

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

- Set up an update server to receive automatic firmware update notifications

  NOTE   This may require a proxy to access HPE public websites on the Internet.

- Update firmware
- Install a license

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 or manual)
  - User definitions (local, LDAP, SNMPv3)
  - Notifications (email, SNMP, syslog)
  - iSCSI settings (host port, network, CHAP)
- Configuring storage settings:
  - Disk group and pool creation
- Provisioning storage:
  - Creating hosts and host groups (naming initiators, assigning initiators to hosts, creating a single host)
  - Creating volumes and attaching them to hosts

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 59. Here is where you begin to use the SMU to monitor, manage, and provision the storage system.

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):
  - 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.

Browser messages warning you about security or privacy concerns due to self-signed or untrusted certificates or invalid certificate authorities are expected if using the system-generated certificates or if the certificate of the CA that signed the device certificate is not in your browser's list of certification authorities. These warnings can be safely bypassed if you trust that you are contacting the correct controller within your network. Depending on the browser and its settings, once you navigate through the browser warning, a security exception may be created and the warning would no longer appear. Your browser address bar may still indicate the connection is not trusted or not secure in this situation. Installing the certificates of the chain of trust CAs in your browser should resolve the browser warnings and indicate that the connection is secure.

Areas of the interface

The main areas of the SMU 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" on the next page.

Table 1 Areas of the SMU interface

Banner:	Product name		⑦ Help	"Setting the date and time" on page 86	"User settings" on page 83	**Log Out**
Menu pane:	"Dashboard" on page 59	**Management pane**				
	"Provisioning" on page 65					
	"Settings" on page 79					
	"Maintenance" on page 93					

Icons in the interface

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

Table 2 Icons in the interface

Icon	Name	Use
⊗	Abort/Cancel	Aborts or cancels an operation.
⊞	Applications	Indicates that the Applications menu is selected.
✓	Apply	Applies an edited operation or selection.
✕	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.
◎	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.
⑦	Information	Opens a small window that defines or provides more information about a feature or option.
ⓘ	Informational	Indicates that an alert or event is informational.
▭	Initiator	Identifies an initiator.
⚲	Maintenance	Indicates that the Maintenance menu is selected.
⚲	Maintenance tasks	Indicates that either maintenance needs to be performed or has already been performed to the specified item.
⊕	Primary replication volume	Identifies the primary replication volume.
⊜	Provisioning	Indicates that the Provisioning menu is selected.

Table 2 Icons in the interface (continued)

Icon	Name	Use
▶	Resume	Resumes a suspended operation.
🗓	Schedule	Indicates that a specified task will take place at defined times.
⤵	Secondary replication volume	Identifies the secondary replication volume.
🔒	Secured	Indicates that the system is secured using FDE.
⚙	Settings	Indicates that the Settings menu is selected.
〉 or 〈	Slide-over arrows	Opens or closes a panel that contains detailed information about an object.
📷	Snapshot	Indicates that a snapshot of the volume was created.
⏸	Suspend	Suspends (pauses) an in-progress operation.
🔓	Unsecured	Indicates that the system is not secured using FDE.
🗄	Volume	Identifies the primary volume.
📤	Volume copy	Indicates the volume is being copied to a new volume.
⚠	Warning	Indicates that an event has Warning severity.

Tips for using the SMU

- Do not use the browser's Back, Forward, Reload, or Refresh buttons. The SMU 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 SMU 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 SMU session.
- As you set options in panels, the SMU 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 SMU.

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. The exported CSV file contains all of the content in the graph. To export performance data, select the ⬆ icon.

Size representations

Parameters such as names of users and volumes have a maximum length in bytes. When encoded in UTF-8, a single character can occupy multiple bytes. Typically:

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

Operating systems usually show volume size in base-2. Disks usually show size in base-10. Memory (RAM and ROM) size is always shown in base-2. 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 SMU, 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.

This setting is available in the SMU as of firmware version IN210. For earlier firmware versions, 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)	$1,000^4$

Table 3 Storage size representations in base 2 and base 10 (continued)

Base-2		Base-10	
Unit	Size in bytes	Unit	Size in bytes
PiB (pebibyte)	$1,024^5$	PB (petabyte)	$1,000^5$
EiB (exbibyte)	$1,024^6$	EB (exabyte)	$1,000^6$

The locale setting determines the character used for the decimal (radix) point, as shown below.

Table 4 Decimal (radix) point character by locale

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

Signing in and signing out

Multiple users can be signed in to each controller simultaneously.

For each active SMU session, an identifier is stored in the browser. All instances of Firefox, Chrome, Edge, and Safari share the same SMU 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 12 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 SMU 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 81 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 64 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. To display the interface in a language different than the one configured for the user, select the language from the user-language list. Language preferences can be configured for the system and for individual users. The default language is English.

4. 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.

Virtual storage

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

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

Some advantages of using virtual storage are:

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

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

- "Automated tiered storage" on page 32
- "SSD read-cache disk groups" on page 21
- "Replication" on page 46
- "Quick rebuild" on page 39
- "Snapshots" on page 36
- "Overcommitting volumes" on page 32
- "Copying volumes or snapshots" on page 37

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

All disks in a disk group must be the same type (SSD, enterprise SAS, or midline SAS). A disk group can contain disks with different capacities, sector formats, and models. If you mix disks with different capacities, the smallest disk determines the logical capacity of all other disks in the disk group regardless of RAID levels except MSA-DP+. 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:

- "MSA-DP+" on page 23
- "Disk sector format" on the next page
- "Virtual disk group configuration" on the next page

Disk sector format

The system supports 512-byte native sector size disks, 512-byte emulated sector size disks, or a mix of these sector formats. The system identifies the sector format used by a disk, disk group, or pool as follows:

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

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

Virtual disk group configuration

A virtual disk group requires the selection of a pool (A or B) along with a specified number of available disks, RAID level, and spare size. If the virtual pool does not exist at the time of adding the disk group, the system will automatically create it. Multiple disk groups (up to 16) can be added to a single virtual pool. A virtual pool may contain from 1 to 16 disk groups. To add disk groups to a pool (**Maintenance** > **Storage**), expand the Disk Groups section for the specified pool (A or B) and then select **Add Disk Group** to access the panel. Panel content is dynamic, displaying options based on protection (RAID) level selected and the available disks.

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

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

Virtual disk groups can be comprised of either all spinning disks or all SSDs. If the system contains only SSDs, it will be treated as an all-flash array. For more information, see "All-flash array" on page 27. Creating disk groups that contain only SSDs in combination with disk groups that contain only spinning disks will allow tiering within a pool. This requires the HPE MSA Advanced Data Services Suite LTU (License To Use) license. For more information, see "Automated tiered storage" on page 32.

If a virtual disk group comprised of spinning disks is created, you can create a read-cache disk group comprised of SSDs without the HPE MSA Advanced Data Services Suite LTU (License To Use) license. A pool may contain either a read-cache disk group or a virtual disk group containing SSDs, but not both. For more information, see "SSD read-cache disk groups" on the facing page.

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

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

Removing disk groups from virtual pools

You can remove one or more disk groups, but not all, from a virtual pool without losing data as long as there is enough space available in the remaining disk groups to move the data into. When a virtual disk group that contains active volume data is removed, that volume data will drain (move) to other disk group members within the pool that have sufficient available capacity. If the pool has enough space to contain the data, the disk group is removed. If the pool does not have enough space to contain the data, the disk group is not removed.

When the last disk group is removed, the pool ceases to exist and will be deleted from the system automatically. Alternatively, the entire pool can be deleted, which automatically deletes all volumes and disk groups residing on that pool.

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

NOTE   If you remove the last disk group in a virtual pool, and the disk group contains at least one volume, a warning prompts you to confirm removing the disk group. If you reply `yes`, the disk group, the pool, and any volumes they contain will be removed; all data in the pool will be lost. If you reply `no`, the disk group, pool, and volumes will remain.

Removing a disk group from a pool

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

SSD read-cache disk groups

A feature of the tiering mechanism is SSD read cache, which uses one read-cache disk group per pool as a read cache for "hot" pages only. SSD read cache is an essentially read-only and capacity-limited tier, and is unrelated to controller read cache.

A read-cache disk group is a special type of disk group that is used to cache pages to improve performance of random reads. Only a single read-cache disk group may exist in a pool. Read cache does not add to the overall capacity of the pool. You can add or remove a read-cache disk group from a pool without any adverse effect on the volumes and their data in the pool, other than to impact random-read performance.

Each read-cache disk group can contain one or two disks, each of which must be an SSD. The maximum usable capacity of a read-cache disk group is the combined capacity of the disks it contains minus 16GiB for metadata, or 3.80TiB, whichever is smaller. A separate copy of the data is also kept in spinning disks. Read cache contents are lost when a controller restart or failover occurs. This does not cause data loss or corruption, as the read cache only duplicates the content that exists in the fault-tolerant disk groups. Taken together, these attributes have several advantages:

- The performance cost of moving data to read-cache is lower than a full migration of data from a lower tier to a higher tier.
- Read cache is not fault tolerant, potentially lowering system cost.
- Controller read cache is effectively extended by two orders of magnitude, or more.

- When a read-cache group consists of one SSD, it automatically uses NRAID. When a read-cache group consists of two SSDs, it automatically uses RAID 0.

You can create a read-cache disk group for a pool if the system is not all-flash array, the system contains available SSDs, and the pool does not contain a disk group comprised of SSDs (known as a Performance tier). A pool cannot contain both SSD read cache and a Performance tier.

To increase the size of read cache within a pool, you must remove the read-cache disk group and then re-add a larger read-cache disk group.

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 MSA-DP+ data protection level, see "MSA-DP+" on the facing page.

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

In the SMU you can create MSA-DP+, RAID-1, RAID-5, RAID-6, and RAID-10 disk groups.

The following tables:

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

Table 5   Example applications and RAID levels

Application	RAID level
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	MSA-DP+

Table 6   RAID level comparison

RAID level	Min. disks	Description	Strengths	Weaknesses
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 1

Table 6   RAID level comparison (continued)

RAID level	Min. disks	Description	Strengths	Weaknesses
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
MSA-DP+	12	Distributed erasure coding with dual disk failure protection using an 8+2 stripe width	Very fast rebuilds, no spare disks (built-in spare capacity), large storage pools, simplified initial deployment and expansion	Requires a minimum of 12 disks for 8+2. This provides 2 disks of spare capacity.

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

RAID level	Number of disks (data and parity)
1	2 total (no parity)
5	3 total (2 data disks, 1 parity disk); 5 total (4 data disks, 1 parity disk); 9 total (8 data disks, 1 parity disk)
6	4 total (2 data disks, 2 parity disks); 6 total (4 data disks, 2 parity disks); 10 total (8 data disks, 2 parity disks)
10	4 to 16 total
MSA-DP+	12 to 128 total

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 with virtual pages. 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 (not SSD 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. For more information, see the Disk group considerations section of the *HPE MSA 2070/2072 Storage Arrays Best practices* white paper.

MSA-DP+

MSA-DP+ 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 MSA-DP+ 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 MSA-DP+ disk group must be the same type (SSD, enterprise SAS, or midline SAS), but can have different capacities. HPE recommends not mixing disks if the ratio of the largest disk to the smallest disk is greater than 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:

- 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.
- When disks have a sizeable difference in capacity, the larger drives are overused in terms of allocated striping causing sequential performance degradation.

Further, HPE recommends to appropriately increase the target space capacity before adding disks. For more information about mixing disks and setting target spare capacity, see the *HPE MSA 2060/2062 Best practices* white paper.

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

Spare disks are not used by MSA-DP+ 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 MSA-DP+ 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. HPE recommends setting the target spare capacity relative to the drive type, capacity, and quantity. See the *HPE MSA 2070/2072 Storage Arrays Best practices* white paper for guidance.

NOTE If a disk fails in an MSA-DP+ 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 *HPE MSA 2070/2072 CLI Reference Guide*. The `spare-capacity` parameter enables setting of the target spare capacity for an MSA-DP+ disk group.

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

Disk group utilities

This section provides information about disk group utilities.

Disk-group expansion

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

When expanding a disk group, all disks in the group must be the same type (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 19.

NOTE To expand a virtual pool, add a disk group as described in "Disk groups" on page 19, or expand an MSA-DP+ virtual disk group.

MSA-DP+ disk groups are expanded when disks are added to the disk group. The system determines how the additional disks are used, either to replenish spare capacity to equal target capacity, to increase usable capacity, or both.

The system 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 MSA-DP+ disk group to expand, access its slide-over panel, select **Expand Disk Group**, and follow the on-screen directions.

Disk-group scrub

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

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

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

The disk-group scrub utility:

- Checks redundancy data (parity) and correct it for protection levels 5, 6, and MSA-DP+.
- Finds, but does not fix, mirror mismatches for protection levels 1 and 10.

  The system reads both copies of mirror data to find any mismatches.
- Finds and fixes media errors for all redundant protection levels.

  Media errors occur when the system cannot read one of the copies of mirror data, due to a disk error such as an Unrecoverable Read Error (URE). (RAID-1 and RAID-10).
  Verify that all blocks are readable (NRAID and RAID-0).

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 95.

SSDs

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 for virtual disk groups in a system that contains only SSDs without the HPE MSA Advanced Data Services Suite LTU (License To Use) license, or in a system that uses SSDs in combination with spinning disks with the HPE MSA Advanced Data Services Suite LTU (License To Use) license. When combined with virtual disk groups that consist of other classes of disks, improved read and write performance is possible through automated tiered storage. Alternatively, you can use one or two SSDs in read-cache disk groups to increase performance of random reads without the HPE MSA Advanced Data Services Suite LTU (License To Use) license. The application workload of a system determines the percentage of SSDs of the total disk capacity that is needed for best performance.

An HPE MSA Advanced Data Services Suite LTU (License To Use) is required when mixing HDDs and SSDs within the same system, except when using SSDs exclusively as SSD read cache. See the following table for examples:

Pool A	Pool B	License required?
HDDs	• HDDs • None	No
SSDs (Flash pool)	• SSDs (Flash pool) • None	
HDDs & SSDs (Read Cache)	• HDDs & SSDs (Read Cache) • HDDs • None	
SSD (Flash pool)	HDDs	Yes
HDD & SSD (Auto-tiering)	• Any configuration • None	

For more information, see:

- "SSD read-cache disk groups" on page 21
- "Gauging the percentage of life remaining for SSDs" below
- "All-flash array" on the facing page
- SSD management
- "SSD read cache" on page 28

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 SMU, 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.

Viewing I/O workload activity

The SMU provides the I/O workload graph to help you configure the target SSD size to get the best performance for your workload.

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

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

Calculations are based on user-specified settings and up to 8 days of usage data captured by the system.

NOTE The suggested capacities might not apply to primarily sequential streaming workloads.

Access the I/O Workload graph from the Dashboard:

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

You can set the following options:

- Values: Select whether to base the calculations on the peak values saved in the usage data or the average values.

  For calculations, the pool is divided into equal "bins" of LBAs. Each sample contains readings for all bins. Multiple samples are taken per day. To calculate *average*, the sum of the readings of a bin are divided by the number of samples. To calculate *peak*, the largest bin value from the collection of samples is taken. This process leaves one value for each bin, whether average or peak is selected. From there, workload calculations are made using the bin numbers as input.

- Show: Select whether to limit the data used for calculations to read I/Os only, write I/Os only, or the combined total of read and write I/Os.

- Workload: Select from one to three workload calculations to display. The default calculations are based on low, mid, and high percentages of capacity: 50%, 80%, and 100%. In place of 50%, you can enter a custom percentage, which must be a whole number.

Reading the graph

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

- When graphed elements are above the SSD capacity line (or if there are no SSDs), data is spread over more capacity in the total system than could be serviced by the SSD capacity. The graph can give you a target SSD size to consider using.

- When graphed elements are below the SSD capacity line, there is adequate SSD capacity for hot data and you're receiving good value from your SSDs.

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

Cumulative heat map data can also be useful in support case analysis. For information, see "Downloading system heat map data" on page 124.

For more information, see:

- "SSD read-cache disk groups" on page 21
- "Gauging the percentage of life remaining for SSDs" on the previous page
- "All-flash array" below
- SSD management
- "SSD read cache" on the next page

All-flash array

An all-flash array is a system whose disk groups that consist of only SSDs; they have one tier that consists solely of SSDs and do not require an HPE MSA Advanced Data Services Suite LTU (License To Use) license. If a system includes disk groups with spinning disks, those disk groups must be removed before the system can be used as an all-flash array. Once the system is established as an all-flash array, only disk groups containing SSDs can be added to the system.

If you are using SSDs and spinning disks and the first disk group is provisioned with spinning disks, then the system will not be treated as an all-flash array. For information about the rules for using SSDs and spinning disks, see "SSDs" on page 25.

SSD read cache

See for information.

Spares

Spare disks are unused disks in your system that automatically replace a failed disk, restoring fault tolerance to disk groups in the system. Designate spares from the **Maintenance** > **Storage** panel. Types of spares include:

- 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. This option is enabled by default.

NOTE MSA-DP+ disk groups do not use global spares or dynamic spares. For information on how MSA-DP+ disk groups manage sparing, see .

The system automatically reconstructs a fault-tolerant disk group (RAID 1, 5, 6, 10) when one or more of its disks fails and a compatible spare disk is available. A disk is compatible if it has enough capacity to replace the failed disk and is the same type (SSD, enterprise SAS, or midline SAS). It is not advisable to mix 10k and 15k disks in a single disk group or pool. If the disks in the system are FDE-capable and the system is secure, spares must also be FDE-capable.

NOTE The system can automatically reconstruct a disk group if it is online and has not suffered more concurrent drive losses than its RAID level can protect against.

When a disk fails, the system 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.

In the SMU 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 RAID 5, RAID 10, or RAID 1 disk group will have 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. A RAID 6 or MSA-DP+ disk group will have critical status until one disk is reconstructed, at which time the disk group will have degraded status; then when the second disk is reconstructed, 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 .

Pools

A pool is an aggregation of one or more disk groups that serves as a container for volumes. A disk group is a group of disks of the same type, using a specific RAID level. For virtual pools, when volumes are added to a pool, the data is distributed across the pool's disk groups.

If the owning controller goes offline, 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.

Virtual pools and disk groups

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

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

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

See "SSD read-cache disk groups" on page 21 for information.

Resolving a pool conflict caused by inserting a foreign disk group

If you insert a virtual disk group, a disk that contains a virtual disk group, or a disk that was actively part of a virtual disk group from one system into another system, the latter system will attempt to create a virtual pool for that disk group. If that system already has a virtual pool with the same name, the pool for the inserted disk group will go offline. For example, if `NewSystem` has pool A and you insert a disk group that came from `OldSystem`'s pool A, the disk group imported from `OldSystem`'s pool A will be offline.

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

IMPORTANT If you are unable to find a pool with a duplicate name, or are unsure of how to safely proceed, download logs from the system and contact HPE Support for assistance with offline disk groups.

Changing pool settings

Each virtual pool has three thresholds for page allocation as a percentage of pool capacity. You can set the low and middle thresholds. The high threshold is automatically calculated.

You can view and change settings that govern the operation of each virtual pool from the pools panel (**Maintenance** > **Storage**): To see information about disk groups in a pool, expand the row. To change pool settings, select the 🖉 icon in the pool row. Options include:

- Low Threshold. When this percentage of virtual pool capacity has been used, the system will generate an alert and informational event 462 to notify the administrator. This value must be less than or equal to the Middle Threshold value. The default is 50%.
- Middle Threshold. When this percentage of virtual pool capacity has been used, the system will generate an alert and event 462 to notify the administrator to add capacity to the pool. This value must be between the Low Threshold and High Threshold values. The default is 75%. If the pool is not overcommitted, the alert will have `Informational` severity. If the pool is overcommitted, the alert will have `Warning` severity.
- High Threshold. When this percentage of virtual pool capacity has been used, the system will generate an alert and event 462 to alert the administrator to add capacity to the pool. This value is automatically calculated as the available capacity of the pool minus 400GiB — or if that result is below the mid threshold, then mid threshold capacity plus 1%. If the pool is not

overcommitted, the alert will have `Informational` severity. If the pool is overcommitted, the alert will have `Warning` severity and the system will use write-through cache mode until virtual pool usage drops back below this threshold.

- Pool Overcommit. This check box controls whether overcommitting is enabled, and whether storage-pool capacity may exceed the physical capacity of disks in the system. For information about overcommitting, see "Overcommitting volumes" on page 32. This option is enabled by default.

NOTE For more information about events, see the Event History panel (**Maintenance** > **Support** > Event History) or see the Event Descriptions Reference Guide.

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

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

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

Volumes

A volume is a logical subdivision of a virtual 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 35.

For virtual pools, when volumes are added to a pool the data is distributed across the pool's disk groups.

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

NOTE About volume groups and replication:

HPE recommends using the SMU to configure replication. Instead of using the CLI to create a volume group, use the SMU to select individual volumes to be included in a SMU-created volume group. If the selected volumes are already members of a CLI-created volume group, the existing volume group will be deleted and replaced with a SMU-created volume group.

If you do use the CLI to create a volume group, use the `create replication-set` CLI command to configure replication for that volume group.

A maximum of 1024 virtual volumes can exist per system.

Volume usage type

When you create a volume, you are required to assign a storage usage type. This metadata identifies how you intend to use the volume. View the usage type for current volumes on the Volumes table (**Provisioning** > **Volumes**).

You can select from the following volume types:

- **Database**: The volume is used for database applications.
- **Virtualization**: The volume is used as a datastore for a virtual machine
- **Email**: The volume is used to hold email or similar content
- **Files**: The volume is used as a file share
- **Other**: The usage type is not described by other choices. Use this option only if none of the other choices meet your needs.

NOTE   The Volumes table might display additional usage values:

- **Unknown**: The volume usage type is unknown, possibly because the volume was created with older firmware.
- **Unspecified**: The volume usage type is not specified because the volume was created by an interface (CLI, API) that did not require a usage value.

These values cannot be assigned when you create or edit volumes. It is recommended that you assign a meaningful usage value instead.

For existing volumes, you can update the usage type by selecting the volume's slide-over to access the **Overview** panel.

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" on the next page
- "Optimizing read-ahead caching " on the next 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 88.

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.

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. This is the default.
*   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.

Overcommitting volumes

Overcommitting is a virtual storage feature that allows a system administrator to allocate storage beyond the capacity of physical storage resources. This allows the host system to operate as though it has more storage available than is present. When physical resources fill up, the administrator can add physical storage by adding additional disk groups or expanding an MSA-DP+ disk group.

Automated tiered storage

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

- Frequently accessed "hot" data can move to disks with higher performance.
- Infrequently accessed "cool" data can move to disks with lower performance and lower costs.
- Each virtual disk group, depending on the type of disks it uses, is automatically assigned to one of the following tiers:
  - Performance: This highest tier uses SSDs, which provide the best performance but also the highest cost. For more information, see "SSDs" on page 25.
  - Standard: This middle tier uses enterprise-class spinning SAS disks, which provide good performance with mid-level cost and capacity.
  - Archive: This lowest tier uses midline spinning SAS disks, which provide the lowest performance with the lowest cost and highest capacity.

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

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

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

Volume tier affinity

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

- No Affinity: This setting uses the highest performing tier available first and uses the other tiers when space is exhausted in the preferred tier; as a result, some tiers may be empty. Volume data is swapped to higher performing tiers based on access frequency and tier space availability. No Affinity is the default setting for the volume tier affinity feature.
- Archive: This setting prioritizes the volume data to the lowest performing tier available. Volume data can move to higher performing tiers based on frequency of access and available space in the tiers.
- Performance: This setting prioritizes writing the volume data to the higher performing tiers more aggressively after first allocation. Performance affinity volume data will swap into higher tiers based upon frequency of access or when space is made available.

All affinity settings are subordinate to the automated tiering engine. For example, the automated tiering engine will use traditional HDDs for sequential workloads and sequential reads will not influence page ranking. For more information on automated tiering, see the *HPE MSA Gen7 Virtual Storage Technical Reference Guide* white paper.

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

The controllers automatically discover initiators that have sent a SCSI `INQUIRY` command or `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 SMU, you must assign a nickname to an initiator in order for it to be added to a host. An initiator can be a member of only one host. A host can be a member of only one group. A host cannot have the same name as another host, but can have the same name as any initiator. A host group cannot have the same name as another host group, but can have the same name as any host. A maximum of 32 host groups can exist. Once you have created a host, you can edit the profile specific to the operating system for that initiator.

CHAP

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

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

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

- Define CHAP records in the storage system.

- Enable CHAP on the storage system (during onboarding or from the **Settings** > **iSCSI** > **Configuration** panel). Note that this applies to all iSCSI hosts, in order to avoid security exposures. Any current host connections will be terminated when CHAP is enabled and will need to be re-established using a CHAP login.

- Define a CHAP record for the host iSCSI initiator on the host.

- Establish a new connection to the storage system using CHAP. The host should be displayable by the storage system, as well as the ports through which connections were made.

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

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

To use CHAP between peers in a replication set, see "CHAP and replication" on page 53.

Host ports

MSA 2070/2072 controller enclosures support FC, iSCSI and SAS host interface protocols. FC and SAS controller host-port settings are not configurable in the SMU.

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

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

Interrelated information about host ports is available from the following sources:

- "Configuring iSCSI host port settings " on page 90.
- For information about setting host parameters, see the CLI Reference Guide.
- For information about host interface protocols, see the Installation Guide.
- For information about supported host port configurations for MSA 2070, see https://www.hpe.com/support/MSA2070QuickSpecs.
- For information about supported host port configurations for MSA 2072, see https://www.hpe.com/support/MSA2072QuickSpecs.

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 69.

When an attachment is created, the system automatically assigns a unique LUN to the volume, sets default permission access to read-write, and sets port access to all ports. After an attachment is created, you can change the LUN, port access, and access permissions. Both controllers share a set of LUNs, and any available LUN can be assigned to a volume.

The storage system uses Unified LUN Presentation (ULP), which can expose all LUNs through all host ports on both controllers. The interconnect information is managed by the controller firmware. ULP appears to the host as an active-active storage system where the host can choose any available path to access a LUN regardless of which controller owns the storage pool the volume resides on. With ULP, the controllers' operating/redundancy mode is shown as Active-Active ULP. ULP uses the T10 Technical Committee of INCITS Asymmetric Logical Unit Access (ALUA) extensions, in SPC-3, to negotiate paths with aware host systems. Unaware host systems see all paths as being equal.

NOTE   LUN 0 is not used for SAS hosts.

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

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

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**)

Snapshots

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

A base of 64 snapshots is included with your system without an additional license. With the HPE MSA Advanced Data Services Suite LTU (License To Use) license, you can create up to the licensed number of snapshots. To view the maximum number of snapshots for your system, see .

When you reach the maximum base number of snapshots for your system, before you can create a new snapshot, you must either purchase or install the HPE MSA Advanced Data Services Suite LTU (License To Use) license that increases the maximum number of snapshots or delete an existing snapshot.

The system treats a snapshot like any other volume:

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

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

Creating snapshots, which uses the redirect-on-write method, is a fast and efficient process that merely consists of pointing to the same data to which the source volume or source snapshot points. (A snapshot takes up no space unless it is directly modified, or the source volume or source snapshot to which the snapshot refers to is modified.) Space does not have to be

reserved for snapshots because all space in the pool is available for them, based upon the limit policy that is set. It is easy to take snapshots of snapshots and use them in the same way that you would use any volume. Since snapshots have the same structure as volumes, the system treats them the same way.

For more information, see:

- "Snapshot trees" below
- "Rollback and reset snapshot features" below

Snapshot trees

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

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

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

Rollback and reset snapshot features

With the rollback feature, if the contents of the selected snapshot have changed since it was created, the modified contents will overwrite those of the source volume or snapshot during a rollback. Since snapshots are copies of a point in time, a modified snapshot cannot be reverted. If you want a snapshot to provide the capability to "revert" the contents of the source volume or snapshot to when the snapshot was created, create a snapshot for this purpose and archive it so you do not change the contents. For more information, see "Rolling back volumes" on page 70.

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

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

You can either reset a snapshot immediately, or schedule the reset. For more information, see "Resetting snapshots" on page 71.

Copying volumes or snapshots

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

The volume copy feature provides the following benefits:

- Additional data protection: An independent copy of a volume provides additional data protection against a complete source volume failure. If the source volume fails, the copy can be used to restore the volume to the point in time when the

copy was created.

- Non-disruptive use of production data: With a mounted independent copy of the volume, resource contention and the potential performance impact on production volumes is mitigated. Data blocks between the source and the copied volumes are independent (versus shared with snapshots) so that I/O is to each set of blocks respectively. Application I/O transactions are not competing with each other when accessing the same data blocks.

The copy operation is performed directly from the source. This source data may change if modified data is to be included in the copy and there are snapshots attached and in use.

You must detach the volume before copying it. It is highly recommended to unmount it on the host first, so that the host can flush its cache, before detaching it.

The volume will not be available for read or write access until the copy is complete, at which time you can reattach the volume.

Reconstruction

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

If sufficient spare capacity is not available, reconstruction does not start automatically. For RAID levels other than MSA-DP+, 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 28).

For virtual storage, reconstruction of all disk groups uses a quick-rebuild feature. For information about quick rebuild, see "Quick rebuild" on the facing page.

For descriptions of LED states, such as for disk failure and reconstruction, see the *HPE MSA 2070/2072 Installation 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.

MSA-DP+ reconstruction

Reconstruction of an MSA-DP+ 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. MSA-DP+ 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 MSA-DP+ 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 MSA-DP+ 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, MSA-DP+ might take longer to reconstruct to fault-tolerant state than a critical RAID-6 running two-disk reconstruct. However, the first phase reconstruction of MSA-DP+, 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 MSA-DP+ 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 MSA-DP+ disk group, and never to let spare space run out. However, if spare space is lost, the `REFT` utility attempts to give the MSA-DP+ disk group the best redundancy across the whole disk group.

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

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

Quick rebuild

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

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

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

Within a few minutes after a quick rebuild completes, a scrub starts on the disk group.

For more information about disk-group reconstruction, see .

Updating firmware

Controller modules, expansion modules, and disk drives contain firmware. Users must have a `manage` role to update the disk or system firmware. The SMU provides options for you to update system firmware, disk firmware, and to use an update server (**Maintenance** > **Firmware**). For information on these options, see:

-
-
-

For more information, see .

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 SFTP/FTP 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. Also, do not attempt to gather performance metrics while performing a log saving operation.

Alternative methods for obtaining log data are to use the **Collect Logs** action (**Maintenance** > **Support**) or the `get logs` command in the SFTP/FTP 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 SFTP/FTP interface, see "Using SFTP/FTP" on page 116.

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 SMU (**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 SMU, see "Managing local users" on page 83. 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 SMU (**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 57.
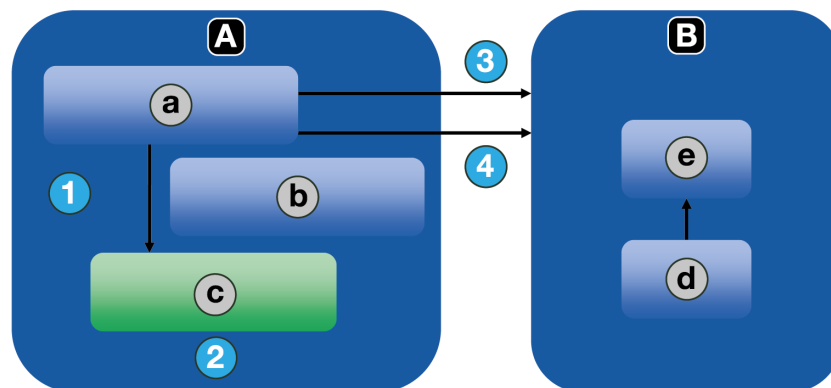
Device		Element		Process	
A	Storage system	a	Session Manager	1	Create session with userid, inherit role from group
B	LDAP server	b	Group: adminusers Role: manage	2	Username will only exist as part of session
		c	Session User: joe Role: manage	3	Authenticate: user
		d	User: joe	4	Check group membership
		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 SMU include the locale, 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 connection. The SMU 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 SMU 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 SMU 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 language, 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 Active Directory 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.

Peer connections

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

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

The prerequisites to create a peer connection are:

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

You can create a maximum of four peer connections per storage system for MSA 2070/2072 Storage systems. However, only one peer connection is allowed to a particular remote system. Attempting to create a second peer connection to the same system will fail.

While creating the peer connection, the local system receives information about all host ports on the remote system as well as the remote system's licensing and host port health. It also links host ports of the selected host port type on the local system to those reachable on the remote system, so all ports of that type are available as part of the peer connection. Once created, the peer connection exists on both the local and remote systems.

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

To create a peer connection, create a replication set by selecting **Provisioning** > **Volumes** > **Data Protection** > **Add Data Protection** > **Remote Replication**. Select one or more volumes to add to the replication set, then follow the on-screen

directions to establish a peer connection between the primary and secondary systems. Directions include steps to enter the secondary system's port address, the connection name, and the username and password of a user with the `manage` role on the remote system. For more information, see "Add data protection" on page 71.

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

For more information about replication, see "Replication" below. For more information about how CHAP interacts with replication, see "CHAP and replication" on page 53. For information about modifying a peer connection, see "Peer connection settings" on page 91.

Replication

Replicating virtual volumes

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

For more information about replication for virtual storage, see "CHAP and replication" on page 53.

Replication prerequisites

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

To verify that a host port address is available before creating a peer connection in the SMU, specify a peer system IP address and then select **Query Peer Connection** (**Settings** > **Peer Connections**). Alternatively, use the `query peer-connection` CLI command. This command provides information about the remote system, such as inter-connectivity between the two systems, licensing, and pool configuration. For more information on this command, see the CLI Reference Guide. For more information on peer connections, see "Peer connections" on the previous page, "Peer connection settings" on page 91, and "Deleting a peer connection" on page 92.

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

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

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

While you cannot modify the secondary volume, you can create a snapshot of the secondary volume that you can map, mount, roll back, and otherwise treat like any volume or snapshot. You can regularly take snapshots to maintain a history of the replications for backup or archiving, or enable snapshot history for the replication set. These snapshots also can be used in disaster recovery. For more information on replication sets, see "Add data protection" on page 71, "Modifying a replication set" on page 72, and "Deleting a replication set" on page 73.

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

Replication process

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

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

The progress and status of the initial and subsequent replications are tracked and displayed. The timestamps for replication reflect the time zones of the respective systems. When viewed on a secondary system in a different time zone, for example, replication information will reflect the time zone of the secondary system. For more information on replicating, see "Aborting a replication set" on page 75, "Initiating or scheduling a replication" on page 73, "Resuming a replication" on page 75, and "Suspending a replication" on page 74.

You can initiate a replication manually or by using a schedule. When creating a schedule for a replication set, you cannot specify for replication to occur more frequently than once per 30 minutes. For more information on scheduling replication, see "Initiating or scheduling a replication" on page 73.

Initial replication

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

Item	Description	Item	Description
A	Replication process - user view	c	Peer connection
B	Replication process - internal view	d	Primary volume
a	Primary system	e	Secondary volume
b	Secondary system		

Figure 2   Replication process for initial replication

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

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

Subsequent replications

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

Item	Description	Item	Description
A	Replication process - user view	c	Peer connection
B	Replication process - internal view	d	Primary volume
a	Primary system	e	Secondary volume
b	Secondary system		
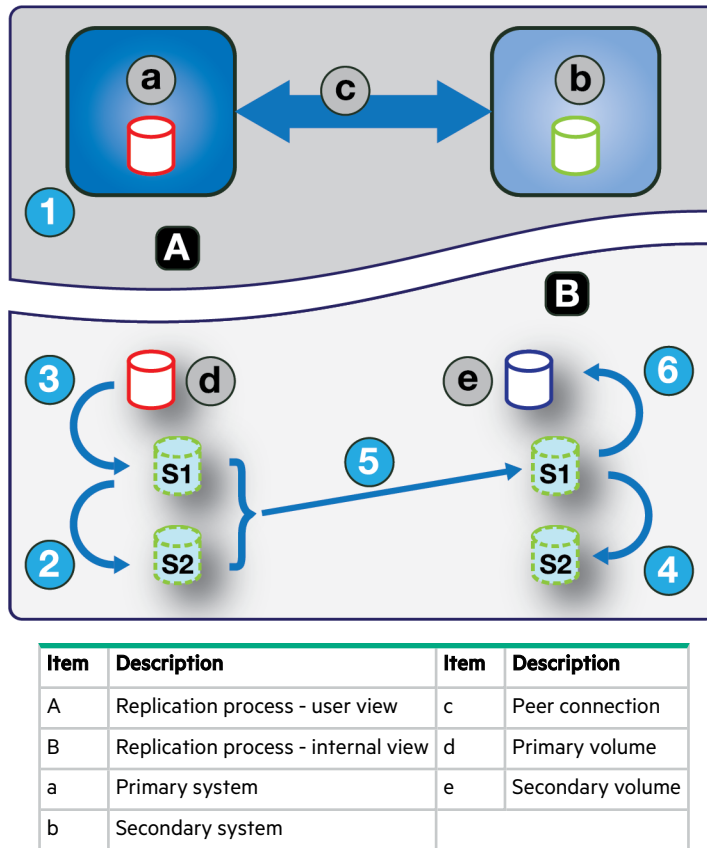
Figure 3   Replication process for replications subsequent to the initial replication

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

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

Internal snapshots

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

Even though the internal snapshots are hidden from the user, they do consume snapshot space (and thus pool space) from the virtual pool. If the volume is the base volume for a snapshot tree, the count of maximum snapshots in the snapshot tree

may include the internal snapshots for it even though they are not listed. Internal snapshots and internal volume groups count against system limits, but do not display and do not count against license limits.

Creating a virtual pool for replication

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

Setting up snapshot space management in the context of replication

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

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

- Make sure that there is enough space in the pool for the base volume and replication-set needs. See guidelines provided in "Creating a virtual pool for replication" above.
- To adjust the snapshot space of the pool, increase the value of the limit parameter of the `set snapshot-space` CLI command. For more information on the `set snapshot-space` CLI command, see the CLI Reference Guide.
- To create more snapshot space, increase the size of the pool, add disk groups to the pool, or expand existing MSA-DP+ groups. It may be necessary to insert additional disks if the system doesn't have any available.

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

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

Limiting the number of snapshots created is the most straightforward approach. You can use the SMU to set the number of snapshots taken by the snapshot history feature. In the **Provisioning** > **Volumes** panel select a volume that is a member of a replication set, select the slide-over arrow to view details, and then select **Replications**; editable parameters are indicated by the 🖉 icon. Reducing the **Snapshot Count** number will save space.

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

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

Replication and empty allocated pages

Deleting data from a volume can result in deallocation of pages on that volume. These pages consume space in the volume and in the pool. Pages deallocated before the initial replication will not be copied to the secondary volume. Pages deallocated since the last replication cause a page consisting of zeroes to be written to the secondary volume during replication. This can result in a difference in the number of allocated pages between the primary and secondary volumes. A virtual storage

background task automatically reclaims pages consisting of all zeroes, eventually freeing up the secondary volume snapshot space that these reclaimed pages consumed. Freeing up this space is not immediate and happens over a period of time.

Disaster recovery

For information about recovering a replication set, see the `recover replication-set` CLI command in the CLI Reference Guide.

Creating a replication set

You can create a replication set, which specifies the components of a replication. The Data Protection Configuration wizard enables you to create replication sets. You can access this wizard from the **Provisioning** > **Volumes** panel. For more information, see .

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

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

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

Primary volumes and volume groups

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

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

Secondary volumes and volume groups

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

Queuing replications

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

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

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

Maintaining replication snapshot history

A replication set can be configured to maintain a replication snapshot history. As part of handling a replication, the replication set will automatically take a snapshot of the secondary and/or primary volume, thereby creating a history of data that has been replicated over time. This feature can be enabled for a secondary volume or for a primary volume and its secondary volume. When this feature is enabled:

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

- For a secondary volume, when a replication successfully completes it will create a snapshot of the data image just transferred to the secondary volume. (This is in contrast to the primary volume snapshot, which is created before the sync.) If replication does not complete, a snapshot of the secondary volume will not be created.

- You can set the number of snapshots to retain from 1 to 16, referred to as the snapshot retention count. This setting applies to management of snapshots for both the primary and secondary volume and can be changed at any time. Its value must be greater than the number of existing history snapshots in the replication set, regardless of whether snapshot history is enabled. If you select a snapshot retention count value that is less than the current number of snapshots, an error message displays. Thus, you must manually delete the excess snapshots before reducing the snapshot count setting. When the snapshot count is exceeded, the oldest unmapped snapshot will be discarded automatically.

- You set the basename of the snapshot. The snapshots are named <basename>_<nnnnnn> where <nnnnnn> starts at 000001 and increments for each subsequent snapshot. If primary volume snapshots are enabled, snapshots with the same name will exist on the primary and secondary systems. The snapshot number is incremented each time a replication is requested, whether or not the replication completes — for example, if the replication was queued and subsequently removed from the queue.

- If the replication set is deleted, any existing snapshots automatically created by snapshot history rules will not be deleted. You will be able to manage those snapshots like any other snapshots.

- If you begin keeping a snapshot history after a number of replications have occurred, the snapshot number in the name will reflect the total number of replications that have occurred.

- Manually creating a snapshot will not increase the snapshot count associated with the snapshot history. Manually created snapshots are not managed by the snapshot history feature. The snapshot history feature generates a new name for the snapshot that it intends to create. If a volume of that name already exists, the snapshot history feature will not overwrite that volume. Snapshot numbering will continue to increment, so the next time the snapshot history feature runs, the new snapshot name will not conflict with that existing volume name.

- The snapshot basename and snapshot retention count settings only take effect when snapshot history is set to secondary or both, although these settings can be changed at any time.

- A mapped snapshot history snapshot will not be deleted until after it is unmapped.

- A snapshot created by this feature is counted against the system-wide maximum snapshots limit, with the following result:
    - If the snapshot count is reached before the system limit, then the snapshot history is unchanged.
    - If the system limit is reached before the snapshot count, then the snapshot history stops adding or updating snapshots.

- The snapshot space management feature, accessible only through the CLI, enables users to monitor and control the amount of space that snapshots can consume in a pool. In addition to configuring a snapshot space limit, you can also specify a limit policy to enact when the snapshot space reaches the configured limit. An event is logged, and if the policy is set to delete, automatic deletion of snapshots occurs. If automatic deletion is triggered, snapshots are deleted according to their configured retention priority.

- You can set the retention priority for snapshots to the following. In a snapshot tree, only leaf snapshots can be deleted automatically.

- never-delete. Snapshots will never be deleted automatically to make space. The oldest snapshot in the snapshot history will only be deleted once the snapshot count has been exceeded. This is the default.
- high. If snapshot space is exhausted, retained snapshots can be deleted after all eligible (system-wide) medium-priority snapshots have been deleted.
- medium. If snapshot space is exhausted, retained snapshots can be deleted after all eligible (system-wide) low-priority snapshots have been deleted, regardless of retention schedule.
- low. If snapshot space is exhausted, retained snapshots can be deleted at any time, regardless of retention schedule.

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

CHAP and replication

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

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

If the two systems have CHAP records for each other with the same secret, they can perform all replication operations whether or not CHAP is enabled on either system. In other words, even if CHAP is not enabled on either system, or if it is enabled on only one system or both systems, either system can work with peer connections, replication sets, and replications.

If you want to use Challenge Handshake Authentication Protocol (CHAP) for the iSCSI connection between peer systems, see the following procedure to configure CHAP. In a peer connection, both systems will alternately act as an originator (initiator) and recipient (target) of a login request. Peer connections support one-way CHAP only.

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

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

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

6. Wait one minute for the commands in step 3 through step 5 to complete before attempting to use the peer connection.
7. Run `query peer-connection` from the local system and then from the remote system to ensure communication can be initiated from either system.

- If both succeed, you can create, set, or perform replication on that peer connection.
- If either fails, it is likely that you must fix a CHAP configuration issue and then repeat step 3 through step 7 as appropriate. If you need to modify a CHAP record, use the `set chap-record` command.

Full disk encryption

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

The ability to secure a disk and system relies on passphrases and lock keys. A passphrase is a user-created password that allows users to manage lock keys. A passphrase can contain up 32 characters and is stored in non-volatile memory. 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.

The lock key 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 and enclosures in the storage system. If both Storage Controllers are online and can communicate with both expansion modules in each connected enclosure, a rescan also reassigns enclosure IDs to follow the enclosure cabling order of controller A. For additional cabling information, see your product's Installation Guide.

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 virtual pool.

The following situations cause a disk to become `LEFTOVR`.

- The disks' timestamps do not match so the system designates members having an older timestamp as leftovers.

- A disk is not detected during a rescan, then is subsequently detected.

- A disk in a disk group is logically or physically removed from the system, and is later returned after the system has noted its removal.

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

- The disk'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** MSA-DP+ 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.

> **NOTE** MSA 2070/2072 controller enclosures do not support transporting cache from one controller module to another.

The default caching mode when a system is operating with a single controller for a volume is write through, as opposed to write back. 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 or the enclosure of the target volume is powered off without a proper shutdown. Data remains in the controller cache and associated volumes will be missing that data on the disk.

If the controller can be brought back online long enough to perform a proper shutdown and the disk group is online, the controller should be able to write its cache to disk without causing data loss.

If the controller cannot be brought back online long enough to write its cache data to disk, contact technical support.

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 31. 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 88.

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 89.

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 16.

Table 8 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 59.

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 SMU, 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 5MB of data. The system will maintain 5MB of audit log data per controller and the log will wrap once the 5MB 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.

System metrics

There are two types of metrics: historical and dynamic. The sampling interval for a dynamic metric is 5 seconds while that of a historical metric is 5 minutes. The system samples every metric at a dynamic sampling interval. The last few minutes of samples are kept in cache and cache is flushed automatically at historical sampling intervals with calculated maximum, minimum, and average data points retained in the historical database. Only a few dynamically sampled user-selected data

series are retained. For such user selected data series, several hours of data points are available. The Performance panel provides charting and data comparison capabilities for metrics that are maintained in the system. For more information, see the topic about the "Performance panel" on page 61.

Effect of failover on metrics

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

If controller A goes offline, all current metrics samples are lost. Controller B becomes active, creates an empty metrics database on the memory card, and begins storing and persisting samples to cover for controller A.

When controller A is back online, controller B clones its current database to controller A. Controller A then initializes its metrics database with the contents of those files.

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 next page
- Changes in performance indicators in the "Performance panel" on page 61
- System activity in the "Activity panel" on page 64

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 14.

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 **Alerts to Acknowledge** 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**). In any of these views, you can select **Export to CSV** to export (download) the data to a comma-separated values (CSV) file for analysis or processing.

To view information about an alert, select an alert type. 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 comments 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 (up to 256 bytes) in the field and then select ✅ to save or ✕ to cancel. For example, you could acknowledge an alert that a FRU is failing and leave a comment saying "Ordered replacement on Oct 14." 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 group and volume level.

Compact view

The compact view on the dashboard shows a capacity graph for each configured pool. Each capacity graph uses a segmented radial bar. The radial bar segments are color-coded as follows:

Color	Represents
	Pool space that is allocated to volumes in the pool.
	Unallocated pool space that remains available to the system.

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:

- System level counts of volumes, snapshots, and disks that are in use versus system configuration limits.
- For each configured pool, the panel shows:
  - The capacity graph, with a link to view the I/O workload.
  - Counts of volumes and snapshots that are in use in that pool versus system configuration limits.

- A pool allocation graph, with color-coding defined on-screen. Pool allocation is represented as a segmented horizontal bar. Text labels on the pool allocation graph report metrics for volume allocated, snapshot allocated, total committed pool allocation, and unused capacity.
  - **Base Volume Allocated**: The amount of thin-provisioned virtual storage used for volume data. This is represented in GB/TB or GiB/TiB (depending on the base preference setting) and as a percentage of the total physical storage capacity.
  - **Snapshot Allocated**: The amount of thin-provisioned virtual storage allocated for all snapshot data. This is represented in GB/TB or GiB/TiB (depending on the base preference setting) and as a percentage of the total physical storage capacity.
  - **Total Committed**: The total storage capacity assigned to volumes. For example, 10 volumes with a capacity of 1TB will result in 10TB of committed storage space.

    NOTE Ensure that sufficient physical storage capacity is available before committing more storage capacity to a volume. Over-committing more storage capacity than available may result in the following:
    - You will see available capacity on a volume but will be unable to use it.
    - An out-of-space error is returned on attempting to perform a new write to a volume.

  - **Unused**: The remaining storage capacity in the pool for creating new volumes or snapshots. This is represented as the pool capacity minus the committed capacity.
- A tiering capacity graph, with color-coding defined on-screen. Tiering capacity is represented as a segmented horizontal bar. Text labels on the tiering capacity graph report metrics for allocated and available capacity on each available tier.
- Volumes for each pool defined (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, attached hosts, creation date, tree size, and capacity utilization.

NOTE The system requires metadata space for volume writes and snapshot creation. When the metadata value reaches zero, the system cannot perform further volume writes or create new snapshots, 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 snapshots or volumes to prevent stoppage of further snapshot creation and volume writes. Once this is done, the metadata metrics should recalibrate to display updated values.

Performance panel

Use the Performance panel to monitor system performance and statistics.

Compact view

The compact view on the dashboard shows consolidated metrics for up to 8 graphs selected as favorites in the expanded view. By default, 3 graphs are shown: System IOPS (read and write), System Latency (read and write), and System Throughput (read and write).

NOTE Compact view does not display live metrics. If the **Live** option is selected in the expanded view, the graphs in the compact view display the most recent historical data for the last 4-hour period.

Select the ⊙ icon to see the expanded view.

Expanded view

The expanded view shows more detailed information and provides options to view metrics within a select time range, add graphs, configure existing graphs, and to select up to 8 graphs as favorites to show in the compact view. You can also choose to display live data up to the last 15 minutes.

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

You can view and compare data-point values by hovering the cursor over the graph, which shows the values in a pop-up table. To see average and peak values for data in a graph, select the ⊡ icon.

Performance metrics

Table 9   Available performance metrics

Metric	Description	Storage object[1,2]				
		Controller	Host port	Pool	System	Volume
IOPS	**Total**: Sum of read IOPS and write IOPS. **Read**: Number of I/Os per second. **Write**: Number of I/Os second.	✓	✓	CLI	✓	✓
Throughput	**Total**: Sum of read bytes per second and write bytes per second. **Read**: Number of bytes read per second. **Write**: Number of bytes written per second.	✓	✓	CLI	✓	✓
Latency	**Average Response Time**: Average response time of an operation in microseconds. Operations include both reads and writes. **Total**: Sum of read and write average responses in microseconds. **Read Average**: I/O read average response time in microseconds. **Write Average**: I/O write average response time in microseconds. **Max Response Time**: Sum of read maximum response time and write maximum response time. **Total**: Sum of read and write maximum responses in microseconds. **Read Max**: Maximum I/O read response time in microseconds. **Write Max**: Maximum I/O write response time in microseconds.	✓	✓	CLI	✓	✓
Average Queue Depth	**Read**: Average number of pending read operations being serviced since the last sampling time. This value represents periods of activity only and excludes periods of inactivity. **Write**: Average number of pending write operations being serviced since the last sampling time. This value represents periods of activity only and excludes periods of inactivity.		✓	CLI		✓
Cache	**Read Ahead Ops**: Number of times that read ahead pre-fetched data for host reads. **Small Destages**: Number of partial stripe destages. (These tend to be very inefficient compared to full stripe writes.) **Write Cache Percent**: Percentage of write cache currently being used in tenths of a percent. **Write Cache Space**: Current size of write cache in 16KB chunks.	✓		CLI		✓

1-Metrics for the Pool storage object are not graphed in the SMU, but are accessible via the CLI.
2-Metrics for the System storage object are synthesized from data captured by Controller storage objects in the storage system.

Collecting performance data

View performance metrics data for a specified period by selecting an option from the **Data Timeframe** dropdown list. Your selection determines your viewing window.

Live performance metrics reveal the system's current performance within the last 15 minutes and are sampled every 5 seconds. The remaining options let you choose to display the system's historical performance up to the last 8 days and are sampled every 5 minutes. All metrics are retained in memory on controller A regardless of whether you are logged in to the SMU.

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

Controller A retains all of the metrics in memory. If controller A is restarted then all metrics will be lost. If controller B is restarted then metrics will continue to be available.

Displaying performance data

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

> 💡 **TIP**   Selecting an object in the legend of a graph highlights that object's data by muting other objects.

The following options are available when in expanded view:

- Select the **Data Timeframe** dropdown to view metrics within a select time range. **Last 1 day** is the default.

  > **NOTE**   Selecting **Live** refreshes the graphs every 5 seconds while in expanded view and displays up to the last 15 minutes of data in 5 second increments. Viewing live graphs can only be done in expanded view. Graphs shown in compact view reflect the data time frame selected in expanded view *except for the Live option*. If **Live** is selected in expanded view, static graphs covering the most recent historical data for the last 4-hour period display when you return to compact view.

- Select the 🔲 icon to view average and peak data values.
- Select **Add Graph** to build a graph by selecting metrics from a predefined list. Graph settings are unique to each graph. The maximum number of objects that can be included in any single graph is four (volumes, host ports, etc.). Additionally, each object can be charted up to three times in the same graph for a comparison (for example, read vs write vs total operations).
- Select the 🤍 icon to designate performance metric favorites (displayed in compact view). Default favorites include System IOPS, System Throughput, and System Latency (displaying both Read/Write).
- Select the ⌃ or ⌄ icon to move the graph up or down in the display.
- Select the ⚙ icon to access the graph's configuration options.
- Select the 📄 icon to download historical performance metrics in CSV format.
- Select the ☒ icon to remove the graph from view. (The graph can be re-added if needed.)

The system retains historical data for one week, or since the last controller restart if that is more recent. By default, the graphs show the latest 100 data points. If a time range or count of samples returns more than 100 data points, then adjacent data samples will be combined until at most 100 data points for the time range or count exist. These combined data points will then be graphed.

IMPORTANT If a volume graph shows an error condition for deleted volumes:

If a volume that was added to a Performance Metrics graph is deleted, then the graph configuration must be manually edited to remove that volume. Otherwise, an error message stating the need to check and update graph settings will appear when the graph displays.

Activity panel

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

Compact view

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

- In Progress: Activities that have started but are not yet complete. This includes disk-group jobs such as initialization and scrub, and long-running application activities such as replicating volumes. A progress bar shows status using green fill and text reporting percent complete.

NOTE The SMU 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 .

4 Provisioning

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

- "Working with volumes" below
- "Working with hosts" on page 75

Working with volumes

The **Volumes** panel (**Provisioning** > **Volumes**) provides options to create volumes, add data protection to volumes, and to display a list of existing volumes and snapshots on the system. For information about creating volumes, see "Creating volumes" on page 68.

If no pools are configured, the panel displays a prompt for you to add storage. If storage is configured the top portion of the panel displays a color-coded graph for each pool that shows the amount of space on the system that is allocated to the base volume, allocated to snapshots, the total committed to volumes (reserved for the system but not written to), and unused space. Each named volume is shown in the Volumes table with columns for type, pool, size, attached hosts, capacity, and a slide-over for access to additional detail and actions.

Volumes table

To perform an action on an existing volume or snapshot, select one or more volumes or snapshots and select an option from the dropdown list:

- "Deleting volumes and snapshots" on page 69
- "Attaching volumes to hosts" on page 69
- "Detaching volumes from hosts" on page 69
- "Expanding volumes" on page 70
- "Add data protection" on page 71
- "Rolling back volumes" on page 70
- "Resetting snapshots" on page 71
- "Copying volumes or snapshots" on page 71

Other actions to perform on an existing volume or snapshot include:

- "Modifying volumes" on page 68
- "Aborting a volume copy" on page 71
- "Creating snapshots" on page 70

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

- The **Overview** tab shows volume capacity usage, volume copy activity, and lists all of the volume settings. Select the ✎ icon to make changes to these settings. Select **Expand Volume** to expand the size of the volume. Select **Copy Volume** to copy the volume. Follow the on-screen directions for more details. For more information, see:
  - "Copying volumes or snapshots" on page 71
  - "Cache optimization mode" on page 32

-
-
-
- The Snapshot tab displays the snapshots associated with the volume, along with any associated schedules. Select the ✎ icon to edit snapshot schedules. The **Snapshot Tree** table lists a hierarchical table of the snapshot tree for the selected volume. Select a volume or a snapshot of the volume, then select an option from the dropdown list to perform an action. Select the parent volume or snapshot, then select **Add Data Protection** and follow the on-screen directions to create local snapshots.
  -
  -
  -
  -
- The **Attached Hosts** tab displays a table listing the hosts attached to the selected volume or snapshot and lets you attach the volume or snapshot to a host or host group. Hover over icons in the list to see more information. 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.
- The **Replications** tab displays the replication sets associated with the volume, along with any associated schedules. Select **Add Data Protection** to access the **Data Protection** wizard; select **Start Replication**, **Suspend Replication**, or **Remove Replication** to perform the action; select the ✎ icon to edit replication schedules and replication details. For more information about replications, see "Replication" on page 46 and "Add data protection" on page 71.

For more information about hosts and initiators, see "Initiators, hosts, and host groups" on page 33. For more information about replications, see "Replication" on page 46 and "Add data protection" on page 71.

Data Protection table

Select **Data Protection** in the Volumes table (**Provisioning** > **Volumes**) to display a list of existing volumes on the system that are protected by a replication set or a snapshot. Select **Add Data Protection** to access the **Data Protection** wizard to protect your volumes using remote replication or local snapshots. To perform an action on one or more volumes in the table, select the volumes and then select an option from the dropdown list:

NOTE   Some options require that only one volume is selected.

-
-
-

-
-

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

- The **Overview** tab displays the volume settings. Select the ✎ icon to make changes to these settings. Select **Copy Volume** to copy the volume. Follow the on-screen directions for more details.

NOTE Snapshots cannot be expanded.

- The **Snapshots** tab displays the snapshots associated with the volume, along with any associated schedules. Select the ✎ icon to edit snapshot schedules. The **Snapshot Tree** table lists a hierarchical view of the snapshots for the selected volume. Select a snapshot, then select **Add Data Protection** and follow the on-screen directions to create local snapshots, or select an option from the dropdown list to perform an action:
  -
  -
  -
  -

IMPORTANT When creating a snapshot, you will be prompted to set a snapshot schedule. This is the only place where you can schedule a snapshot using the SMU. If you do not set a snapshot schedule when prompted, you will be unable to do so using the SMU after the snapshot is created. For information about creating a snapshot schedule using the `set schedule` CLI command, see the CLI Reference Guide.

- The **Attached Hosts** tab displays a table listing the hosts attached to the selected volume or snapshot and lets you attach the volume or snapshot to a host or host group. Hover over icons in the list to see more information. 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.
- The **Replications** tab displays the replication set associated with the volume, along with any associated schedules. If there is no existing replication set for the volume, select **Add Data Protection** to access the **Data Protection Configuration** wizard; if there is a replication set for the volume, select **Start Replication**, **Suspend Replication**, or **Remove Replication** to perform the action; select the ✎ icon to edit any replication schedules and replication details. For more information about replications, see "Replication" on page 46 and "Add data protection" on page 71.

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

See the following topics for more information:

- "Snapshots" on page 36
- "Volumes" on page 30
- "Replication" on page 46
- "Initiators, hosts, and host groups" on page 33
- "Copying volumes or snapshots" on page 37
- "Data protection with a single controller" on page 56

Creating volumes

Select **Create Volumes** (**Provisioning** > **Volumes**) to open the **Create Volumes** wizard to add volumes to a pool. The top portion of the panel displays a color-coded graph for each pool on the system. The graph indicates the amount of space on the system that is allocated to volumes; the total volumes committed (reserved for the system but not written to); the space occupied by added volumes; the space required for the new volume you are creating; unused; and the space that is overcommitted if the Pool Overcommit setting is enabled.

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 30
- "Initiators, hosts, and host groups" on page 33

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 134.

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 the volume, copy the volume, edit the volume name, and change the volume's usage type, tier affinity, cache write policy, cache optimization mode, and cache read-ahead size. If the volume is not a snapshot or a secondary volume involved in replication, you can expand the volume size.

For more information, see:

- "Volume cache options" on page 31
- "Volume usage type" on page 30
- "Volume tier affinity" on page 33

⚠️ **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 and snapshots

You can delete volumes and snapshots from the:

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

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

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

Select a volume or snapshot, then 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, snapshots, or both) to delete.
- You cannot delete a volume that is part of a replication set.
- Ensure that hosts are not accessing the volumes or snapshots to be deleted.

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

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

Attaching volumes to hosts

Attach volumes to hosts from 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**)
- Data Protection table (**Provisioning** > **Volumes** > **Data Protection** > slide-over > **Attached Hosts** > **Unmap**)
    Follow the on-screen directions to complete the action.

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

Expanding volumes

You can expand the size of a volume from 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**)
- Data Protection table (**Provisioning** > **Volumes** > **Data Protection** > slide-over > **Overview** > **Expand Volume**)

  Follow the on-screen directions to complete the action.

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

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

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

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

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

Rolling back volumes

You can replace the data of a source volume or snapshot with the data of a snapshot that was created from it by accessing the:

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

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

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

For more information, see .

Creating snapshots

You can create snapshots of a volume by accessing the:

- Volumes table (**Provisioning** > **Volumes**), selecting the volume, and selecting **Add Data Protection** from the dropdown list.

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

Follow the on-screen directions to complete the action. For more information about snapshots, see "Snapshots" on page 36.

Resetting snapshots

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

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

Select the volume and select **Reset Snapshot** from the dropdown list. Follow the on-screen directions to complete the action.

NOTE Any snapshot in a snapshot tree can be reset, but the data source can only be the snapshot's immediate parent. The snapshot's volume characteristics are not changed.

Copying volumes or snapshots

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

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

Select the volume and select **Copy Volume** from the dropdown list. Follow the on-screen directions to complete the action. For more information, see "Copying volumes or snapshots" on page 37.

NOTE You must detach the volume before copying it.

Aborting a volume copy

You can abort a volume copy operation (**Provisioning** > **Volumes**) by selecting the slide-over of the volume being copied. In the **Overview** panel, select the ⊗ icon next to the progress indicator. Follow the prompts to abort the operation.

Add data protection

Selecting **Add Data Protection** opens the **Data Protection Configuration** wizard where you are led through the process of adding data protection to a selected volume by either creating local snapshots or creating a remote replication set. The options that display in the wizard are dependent upon how the wizard is accessed:

- **Provisioning** > **Volumes** > menu list > **Add Data Protection**
- **Provisioning** > **Volumes** > **Data Protection** > **Add Data Protection**
- **Provisioning** > **Volumes** > slide-over > **Snapshots**
- **Provisioning** > **Volumes** > slide-over > **Replications**

Follow the on-screen directions to complete the action.

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

See the following topics for more information:

- "Replication" on page 46
- "Snapshots" on page 36
- "Modifying a replication set" below
- "Deleting a replication set" on the facing page
- "Initiating or scheduling a replication" on the facing page
- "Aborting a replication set" on page 75
- "Suspending a replication" on page 74
- "Resuming a replication" on page 75
- "Managing replication schedules" on page 75

Creating a replication set

Selecting **Add Data Protection** opens the **Data Protection Configuration** wizard where you are led through the process of creating a replication set for a selected volume. The options that display in the wizard are dependent upon how the wizard is accessed:

- **Provisioning** > **Volumes** > menu list > **Add Data Protection**
- **Provisioning** > **Volumes** > **Data Protection** > **Add Data Protection**
- **Provisioning** > **Volumes** > slide-over > **Replications**

Follow the on-screen directions to complete the action.

NOTE About volume groups and replication:

HPE recommends using the SMU to configure replication. Instead of using the CLI to create a volume group, use the SMU to select individual volumes to be included in a SMU-created volume group. If the selected volumes are already members of a CLI-created volume group, the existing volume group will be deleted and replaced with a SMU-created volume group.

If you do use the CLI to create a volume group, use the `create replication-set` CLI command to configure replication for that volume group.

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

Modifying a replication set

You can modify a replication set's name, queue policy, snapshot history settings, and associated schedules (**Provisioning** > **Volumes**) by selecting the slide-over to access the **Replications** panel. Select the 🖉 icon next to the options you want to modify in the Replication Details section.

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

☼ **TIP** Consider the following points before modifying a replication set:

- If you change the Snapshot Base Name while a replication is running, for the current replication it will affect the name of the snapshot on the secondary system. For that replication only, the names of the snapshots on the primary and secondary systems will differ.

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

See the following topics for more information:

- "Queuing replications" on page 51
- "Maintaining replication snapshot history" on page 52

Deleting a replication set

You can delete a replication set (**Provisioning** > **Volumes**) by selecting the slide-over to access the **Replications** panel. Select **Remove Replication**. The replication set is deleted.

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

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

For more information, see:

- "Aborting a replication set" on page 75
- "Suspending a replication" on the next page
- "Maintaining replication snapshot history" on page 52

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

Initiating or scheduling a replication

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

- If a replication is not in progress, the local system begins replicating the contents of the replication set volume to the remote system and a progress bar indicates the status of the replication set.

- If a replication set is already in progress, then the outcome of this replication request depends on the Queue Policy setting specified. For more information on setting the queue policy, see "Queuing replications" on page 51.

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

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

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

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

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

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

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

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

Suspending a replication

You can suspend replication operations for a specified replication set (**Provisioning** > **Volumes**) by selecting the slide-over to access the **Replications** panel. If the replication is in progress, select the ⏸ icon next to the progress indicator. If the replication is in an unsynchronized or ready state, select **Suspend Replication**.

You can suspend replications only from a replication set's primary system.

When you suspend a replication set, all replications in progress are paused and no new replications are allowed to occur. The replication set will stay in a suspended state until you resume it by selecting the ▶ icon, or abort it by selecting the ⊗ icon. Resuming the replication set continues the replications that were in progress and allows new replications to occur. Aborting the replication set cancels the replication. For more information, see "Aborting a replication set" on the facing page or "Resuming a replication" on the facing page.

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

Aborting a replication set

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

You can abort replications only from a replication set's primary system.

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

Resuming a replication

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

You can resume replications only from a replication set's primary system.

When a replication set is suspended, all replications in progress are paused and no new replications are allowed to occur. When you resume replications, all paused replications are resumed and new replications are allowed to occur depending on the queue policy setting. For more information on setting the queue policy, see "Queuing replications" on page 51.

Managing replication schedules

You can modify or delete scheduled replication tasks from the primary system (**Provisioning** > **Volumes**) by selecting the slide-over to access the **Replications** panel. Select the 🖉 icon in the Replication Schedules table and follow the on-screen directions to modify the schedule. Select the 🗑 icon to delete the schedule.

IMPORTANT   This option is unavailable if a replication set schedule is not defined. See "Add data protection" on page 71 for more information.

Working with hosts

The Hosts panel (**Provisioning** > **Hosts**) provides options to create hosts and host groups, display a list of existing hosts, host groups, and initiators that are a part of an existing host or host group, and display a list of all initiators. For more information about creating hosts, see "Creating hosts" 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" on the next page
- "Detaching hosts from volumes " on the next page
- "Removing initiators from a host" on page 77
- "Removing hosts from a host group" on page 77
- "Adding hosts to a host group" on page 77
- "Deleting hosts" on page 77
- "Deleting host groups" on page 77

Other actions to take on this tab include:

- "Renaming hosts " on page 78
- "Changing a host profile" on page 78
- "Renaming host groups" on page 78
- "Renaming initiators" on page 78

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 33.

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.

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

For more information about volumes, see "Volumes" on page 30. For more information about hosts, host groups, and initiators, see "Initiators, hosts, and host groups" on page 33.

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 directions 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 83
- "System settings" on page 85
- "Notification settings" on page 89
- "Configuring iSCSI host port settings " on page 90
- "Peer connection settings" on page 91

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, 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 81
- "Viewing certificate information" on page 82
- "Configuring a proxy server" on page 82

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. Alternatively, each controller can be set to use IPv6 only, which disables all IPv4 communication on that controller. IPv6-only mode supports use cases where IPv4 is forbidden for security reasons.

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.

To use IPv6 only, select the **IPv6 Only Mode** option. This option is disabled by default. Enabling IPv6-only mode on a controller will disable all IPv4 options on the **Settings** > **Network** > **IPv4** tab for that controller and affect the following services: DNS, proxy server, LDAP, SNMPv3, system date and time (if an IPv6 NTP server is used), and notification (email, SNMP, and syslog). Those services will have to be configured with an IPv6 address. If they are already configured with an IPv4 address, the related options in the SMU will be identified with a red asterisk and a tooltip to indicate they must be reconfigured.

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.

The factory-default IP address source is set to DHCP. When DHCP is enabled in the storage system, the following initial values are set and remain set until the system is able to contact a DHCP server for new addresses:

- 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: 10.0.0.0

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

⚠ **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 .

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 44.

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 SMU, CLI, or other interfaces, see "User settings" on page 83.

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

- **HTTP**. Enable the use of HTTP to provide access to the SMU. By default, HTTP is disabled.
- **HTTPS**. Enable the use of HTTPS to provide secure access to the SMU, the web application that is the primary interface for managing the system. By default, HTTPS is enabled.

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

Command Line Interface

- **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**. 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.
- **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, installing a license, and downloading logs. FTP is disabled by default.
- **SFTP**. A secure secondary interface for uploading firmware updates, downloading logs, installing a license, 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. This is disabled by default.
- **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. SLP is disabled by default.

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 the controller, select **Maintenance** > **Hardware** > **Rear View** > **Enclosure Actions** > **Restart/Shutdown System** > **Full Restart MC** and follow the on-screen directions, 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.

For details pertaining to supported certificate criteria, see .

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" below](#)
- ["Managing LDAP users" on the next page](#)
- ["Managing SNMPv3 users" on page 85](#)

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, language and temperature preference, and timeout setting. Standard and monitor users cannot change their access to user interfaces or roles, and they cannot change the settings of other users.

Users having the `manage` or `standard` role can access one or more of the following management interfaces: the SMU, CLI, or SFTP/FTP. Monitor users can only access the SMU 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 64 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 " ' , < > \

- **Language**. Select a display language for the user. The default is English. Installed language sets include Chinese-Simplified, Chinese-Traditional, Dutch, English, French, German, Italian, Japanese, Korean, and Spanish. The locale determines the character used for the decimal (radix) point, as shown in "Size representations" on page 16. The locale setting is determined by the Language setting, which can be accessed by selecting the pencil icon for any user in the table.

- **Temperature Preference**. Select whether to use the Celsius or Fahrenheit scale for display of temperatures. The default is Celsius.

- **Timeout**. Select the amount of time that the user'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 SMU. This is the default.
  - **CLI**. Enables access to the command-line interface. This is the default.
  - **FTP**. Enables access to the FTP interface or the SFTP interface, which can be used instead of the SMU 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 SFTP/FTP, performing file uploads from the SMU (WBI), or using the CLI `restore defaults` command.
  - **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, language and temperature preference, and a timeout setting.

Users logging into the SMU 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 41.

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 SMU. This is the default.
  - **CLI**. Enables access to the command-line interface. This is the default.
  - **FTP**. Enables access to the SFTP interface or the FTP interface, which can be used instead of the SMU to install firmware updates, to install licenses, and to download logs. The SFTP/FTP interface is also used to upload certificates, and download I/O density data and disk performance statistics.
- **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 SFTP/FTP, performing file uploads from the

SMU, or using the `restore defaults` CLI command.

- **Monitor**. Enables the user to view but not change system status and settings.
- **Language**. Select a display language for the user. The default is English. Installed language sets include Chinese-Simplified, Chinese-Traditional, Dutch, English, French, German, Italian, Japanese, Korean, and Spanish. The locale determines the character used for the decimal (radix) point, as shown in "Size representations" on page 16.
- **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 107.

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.

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 64 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. The default is MD5. 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. The default is none. 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.
- **Trap Host Port**. Specify the target port of the host that will receive SNMP traps. The default port is 162.

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.

-
-
-
- (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 SMU 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>`. The banner also displays the configured time zone region, or `GMT` if no region is set.

NOTE   You can also access the **Date and Time** panel by selecting 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 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. An NTP server provides a valid global epoch timestamp that allows multiple storage devices, hosts, log files, and so forth to be synchronized across multiple time zones. 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.

You can also set the system's time zone region. After you do so, all system interfaces and logs will display the local time according to the configured time zone. The system clock will automatically adjust for daylight saving time (DST).

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

Securing the system with FDE

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

The **Security** panel shows whether the system is secured.

To secure the system, select **Secure System** to enter a passphrase to generate the lock-key ID and enable data-at-rest security. If the disks were previously secured then you must use the previous passphrase to resecure the system. If a disk was previously secured with a different lock-key ID, the disk's encrypted data is inaccessible.

IMPORTANT   Be sure to record the passphrase as it cannot be recovered if lost.

After setting the options and selecting **Secure System**, the **System Status** property will show that the system is secured and the **Security** panel will provide options to:

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

- Repurpose secured disks, as described in "Repurposing secured disks" below.
- Repurpose the system. This action will erase all data on secured drives and return its FDE state to unsecured. This action is only enabled when all pools have been removed from the system.

⚠ **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 unlocks the disk 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   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 54.

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" on the next page
- "Setting system disk properties" on the next page
- "Setting system scrub properties" on the next page
- "Setting system managed logs properties" on page 89
- "Setting partner firmware update" on page 89

IMPORTANT Settings that could potentially prevent the system from detecting or recovering from failure:

- Disabling background scrub will prevent automatic detection of disk defects for disk-group members.
  (Enabled by default)
- Disabling SMART (Self-Monitoring Analysis and Reporting Technology) will prevent the system from monitoring disk failures.
  (Enabled by default)
- Disabling dynamic sparing will prevent the system from automatically integrating replacement disks into disk-groups.
  (Enabled by default)
- Disabling auto-stall-recovery will prevent the system from automatically recovering at least one controller to continue operation when a controller stall is preventing I/O.
  (Enabled by default)
- Disabling controller failure (Enabled), fan failure (Disabled), supercap failure (Enabled), power supply failure (Disabled), overtemperature failure (Disabled), and notify controller (Enabled) will prevent automatically changing the controller cache policy from write-back to write-through when those conditions occur.
  (Trigger default settings are noted above)

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 disk spin down (DSD)

See the tool tips in the panel for specific information regarding each option. For more information about dynamic spares, see "Spares" on page 28.

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 enabling the **Disk Group Scrub** option. 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 25.

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 **Disk Group Scrub** is enabled, you can select **Create A Schedule** to schedule scrub operations. When creating a schedule, both a schedule start and schedule stop time are required. Multiple schedules can be created for scrub but be careful 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.

For more information, see the *HPE MSA 2070/2072 Storage Arrays Best practices* white paper.

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

Setting partner firmware update

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

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

Notification settings

The **Notifications** panel (**Settings** > **Notifications**) provides options to send system alert notifications to users through email, SNMP trap hosts, or a remote syslog server. For more information about alerts, see "Alerts panel" on page 59.

💡 **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 *HPE MSA 2070/2072 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 79.

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, and be sure to use different values for the read-community and write-community strings. For details about panel options, see the on-screen tool tips. See "Enabling or disabling system-management services" on page 81 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.

Configuring iSCSI host port settings

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

Reset Host Links. Making a configuration or cabling change on a host might cause the storage system to stop accepting I/O requests from that host. For example, this problem can occur after moving host cables from one HBA to another on the host. To remedy this, you might need to reset controller host links (channels).

IP Address. For IPv4 or IPv6, the port IP address. For corresponding ports in each controller, assign one port to one subnet and the other port to a second subnet. Ensure that each iSCSI host port in the storage system is assigned a different IP address. For example, in a system using IPv4:

- Controller A port 1: 10.10.10.100
- Controller A port 2: 10.11.10.120
- Controller B port 1: 10.10.10.110
- Controller B port 2: 10.11.10.130

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

Netmask. For IPv4, enter the subnet mask for the assigned port IP address. The default is 255.255.255.0

Gateway. For IPv4, enter the gateway IP address for the assigned port IP address. The default is 0.0.0.0.

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

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

Configuring iSCSI CHAP settings

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

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

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

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

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

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

For more information, see:

- "Peer connections" on page 45
- "CHAP and replication" on page 53
- "Initiators, hosts, and host groups" on page 33

Changing iSCSI configuration settings

The **iSCSI** configuration panel (**Settings** > **iSCSI** > **Configuration**) provides options for you to view the system's configuration, or to alter it if the network configuration was modified. Panel options include the ability to:

- Change the IP version. IPv4 uses 32-bit addresses. IPv6 uses 128-bit addresses.
- Set the host port link speed.
- Enable/disable jumbo frames to allow for larger data transfers.
- Enable/disable CHAP authentication.
- Enable/disable iSNS.

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

Peer connection settings

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

Querying peer connections

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

Modifying peer connection settings

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

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

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

If **CHAP** is enabled on one system within a peer connection, be sure that CHAP is configured properly on the corresponding peer system before initiating this operation. For more information about configuring CHAP, see "CHAP and replication" on page 53.

Deleting a peer connection

Selecting the 🗑 icon from the **Current Peer Connections** section of the panel lets you delete a peer connection if there are no replication sets that belong to the peer connection. If there are replications sets that belong to the peer connection, you must delete them before you can delete the peer connection. For more information, see "Deleting a replication set" on page 73.

NOTE If the peer connection is down and there is no communication between the primary and secondary systems, use the `local-only` parameter of the `delete replication-set` CLI command to delete the replication set.

NOTE If CHAP is enabled on one system within a peer connection, be sure that CHAP is configured properly on the corresponding peer system before initiating this operation. For more information about configuring CHAP, see "CHAP and replication" on page 53.

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 95
- "Firmware panel" on page 96
- "About panel" on page 101
- "Support panel" on page 102

Storage panel

The **Storage** panel (**Maintenance** > **Storage**) shows the system's storage configuration, including pools, disk groups, and spares and enables you to change the configuration. To learn about pools and disk groups, see "Pools" on page 28. To learn about spares, see "Spares" on page 28.

A storage system can have one pool per controller module. This panel shows a Pool table for each pool, and a Spares table. In this panel you can:

- View information about a pool
- Change pool settings
- View information about disk groups in a pool
- Add a disk group to a pool
- Rename a disk group
- Delete a disk group from a pool
- Expand an MSA-DP+ disk group
- Scrub a disk group

Viewing information about a pool

If a pool contains at least one disk group, the Pool table shows this basic information:

- Size (total capacity)
- Health
- Available (unused capacity)
- Overcommit size

Expand the pool row to see this additional information:

- Whether overcommit is enabled
- Whether the pool is overcommitted
- Low threshold, medium threshold, and high threshold values
- Pool serial number

Changing pool settings

To change pool settings, in the pool row select the ✎ icon. For more information about each setting, select its ⑦ icon.

Viewing information about a disk group

To see information about disk groups in a pool, in a Pool table expand the Disk Groups row. For each disk group, the Disk Groups table shows this basic information:

- Name
- Level (disk-protection level)
- Type
- 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 (MSA-DP+)
- Actual spare capacity (MSA-DP+)

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>`.

Adding a disk group to a pool

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 directions.

For more information about available protection levels, see "RAID levels" on page 22.

Renaming a disk group

In the **Maintenance** > **Storage** panel, locate the disk group to rename, display its slide-over panel, select the ✏ icon, and follow the on-screen directions.

Deleting a disk group from a pool

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

Expanding an MSA-DP+ 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 MSA-DP+ disk group. Target spare capacity may only be increased when expanding the disk group.

See the *HPE MSA 2070/2072 Storage Arrays Best practices* white paper for more information.

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 .

Hardware panel

The Hardware panel (**Maintenance** > **Hardware**) shows the system's hardware configuration, including attached enclosures and their components.

The panel has three sections:

- The top section shows basic information about each enclosure: ID, Rack number, Rack position, Disk slots (used and total).
- For the selected enclosure, the middle section shows a front view or rear view of the position of components in the enclosure. To the right of the enclosure view, a box shows the health of the selected component and available actions that pertain for that component.
- For the selected enclosure or component, the bottom section shows additional information.

This table lists available actions for a given device.

View	Device	Available actions
Front or Rear	Enclosure	Restart/shutdown system
		Rescan all disks
		Turn on locator LED
Front	Disk (healthy)	Turn on locator LED
	Disk (leftover)	Turn on locator LED
		Clear disk metadata

Rear	Power supply	None
	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, EMP A revision, EMP B revision, EMP A bus ID, EMP B bus ID, EMP A target ID, EMP B target ID
Disk module	Location, Locator LED On/Off button, LED status, Serial number, Vendor, Model, Revision, Description, Usage, Current job, Supports unmap, SMART, R/min (RPM), Size, Sector format, Transfer rate, Single pathed, Recon state, Copyback state, Disk spin down count, Temperature, Status, Power on hours, SSD life remaining, FDE state, FDE lock key. Disk modules reside in disk slots and are accessible from the front of the enclosure.
Power cooling module	Status, Vendor, Model, Serial number, Revision, Location, Part number, Manufacturing date, Manufacturing location. The two power cooling modules (PCM), numbered 1 and 2, reside in the PCM slots accessed from the rear of the enclosure.
Controller module	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.
FC host port	Name, Port type, Status, Topology, Configured speed, Actual speed, Primary loop ID, Target ID, SFP status, Part number, Supported speeds
iSCSI host port	Name, Port type, Status, Gateway, Netmask, MAC address, IP address, IP version, ID, SFP status, Part number, Configured speed, Actual speed, 10G Compliance, Cable length, Cable technology, Ethernet compliance
SAS host port	Name, Port type, Status, Actual speed, Topology, Expected lanes, Active lanes, Disabled lanes, ID
Network port	ID, IPv4 address mode, IPv4 address, IPv4 network mask, IPv4 gateway, MAC address, IPv6 only mode, IPv6 auto config, IPv6 gateway, IPv6 auto address, IPv6 manual address (1 to 4)
Expansion port	Enclosure ID, Controller ID, Name, Status

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 accessible for 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. You can also configure access to an update server.

The system stores two bundles of system 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

- Configure access to an update server

- View update server information

💡 **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 additional 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 version (Storage Controller)

Updating system firmware

Before performing a firmware update, see .

In an MSA 2070/2072, both controllers must run the same firmware version. Storage systems with peer connections should run the same or compatible firmware versions. You can find and install an HPE Smart Component from the website: https://www.hpe.com/storage/msafirmware. Follow the instructions on the HPE website.

NOTE HPE recommends using the HPE Smart Component when updating firmware.

In the **System** tab, the process to update firmware is to install firmware bundles obtained from HPE and then activate a particular bundle. 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. PFU provides for updating expansion module firmware in similar fashion.

> **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.

If the system is connected to the Update Server and an update is available, a link will display to retrieve the latest firmware. See "Using an update server" on page 100 for more information.

For information about supported releases for firmware update, see the MSA Storage Firmware page at https://www.hpe.com/storage/msafirmware.

To install a firmware bundle, follow the on-screen directions and "Best practices for updating firmware" on page 101.

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. Using the recommended HPE Smart Component installs and automatically activates the newly installed firmware. 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 SMU 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 controller enclosure has connected enclosures, allow additional time for each expansion module's enclosure management processor (EMP) to be updated. This typically takes 2.5 minutes for all EMPs in an MSA 2070/2072 disk enclosure.
- 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 expansion module firmware

An expansion enclosure contains two expansion modules. Each expansion module contains an EMP. All modules of the same model should run the same firmware version. If PFU is enabled, the system automatically updates expansion modules from the active bundled firmware when new firmware becomes available. If PFU is enabled, a replacement expansion module will be updated automatically. For more information, see relevant command topics below.

NOTE See the following interrelated commands in the *HPE MSA 2070/2072 CLI Reference Guide:*

- `activate firmware`
- `check firmware-upgrade-health`
- `check update-server`
- `show firmware-bundles`
- `show firmware-update-status`
- `show system`
- `show update-server`
- `show versions`

Updating disk firmware

If the system is connected to the Update Server and an update is available, a link will display to retrieve the latest firmware. See "Using an update server" on the next page for more information.

NOTE HPE only provides the HPE Smart Component when updating disk firmware. To install an HPE Smart Component, follow the instructions on the HPE website at https://www.hpe.com/storage/msafirmware.

Using an update server

Configuring the update server for the storage system enables the system to periodically retrieve a manifest of available firmware, IOM, or disk firmware update components from HPE. If system or disk firmware updates are available, alert notifications are generated. The General section on the **Firmware** panel displays the most recent date and time that the system checked for available manifest updates (**Last Check**) and the most recent status of the updates (**Last Status**).

In order to retrieve the update notifications, a connection to HPE's publicly available manifest is required. This URL points to a manifest file that documents supported upgrade paths. By default, the URL for the update server is https://www.hpe.com/support/MSAmanifest.

NOTE A valid DNS server must be configured for the update server manifest to be located as configured from the factory.

To configure the update server, go to **Maintenance** > **Firmware** > **Update Server**. The configured update server URL displays in the **Resource URL** text box. Select **Set Update Server**. Once access to an update server is configured, the **General Information** section shows the update server's **Resource URL** and the time, date, and status of the last check.

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 to receive the manifest. Only HTTP is supported to the proxy server, though the proxy itself can support HTTPS traffic to the final endpoint.

See "Configuring a proxy server" on page 82 for information about configuring a proxy server and setting proxy parameters.

To manage updates locally using the update server within your firewall, configure the storage system to work with a local copy of the manifest on an internal website:

- Download the publicly available manifest from https://www.hpe.com/support/MSAmanifest and place it on an internal website. Note the manifest's URL.

  NOTE The manifest is a standard JSON (JavaScript Object Notation) file.

- Configure the storage system to access the manifest that was downloaded to the internal website by entering its URL in the **Resource URL** text box and then select **Set Update Server**. This enables the system to generate alerts and send notifications on available updates. Updates display in the **Status** column of the **System** tab or the **Update** column of the **Disks** tabs.

  NOTE The update server retrieves only the manifest to validate and process updates within the MSA controller. It does not send or retrieve any data.

- Download firmware component updates using a system with Internet access from https://www.hpe.com/support.

  NOTE Some components require product entitlement and a valid HPE Account.

- Post the downloaded firmware components to an internal site and edit the manifest JSON file to point to the internal component location. Doing so allows each MSA system to have access via links in the SMU to the firmware updates applicable to that specific system.
- Regularly update the manifest on the internal server with HPE's publicly available manifest to be aware of the most recent update components available.

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 59.
- 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.

  $\triangle$ **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**) provides links to useful websites and 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

- Controller A firmware
- Controller B firmware

Hardware information displays the enclosure number, SKU part number, and SKU serial number for each enclosure (disks excluded).. Expand the table to see the following for each customer FRU in each enclosure:

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

Storage information includes the following for each pool:

- Pool ID
- Total size
- Available size
- Snap size
- Overcommit state
- Disk groups
- 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's licensing serial number and install a license for advanced features
- View the system event history and export (download) the data to a comma-separated values (CSV) file for analysis or processing
- View controller module audit logs

For details, see the on-screen directions.

7 Applications

Use the **Applications** panel to manage custom applications added to the storage system through use of the MC Auxiliary Services (MAS) feature. These applications are also referred to as MAS services or agents.

If no MAS applications are installed, the **Applications** menu option and panel do not appear.

Application information

The **Application Information** tab shows information about each MAS application, including:

- URL
- State
  - `Enabled`: The application is running.
  - `Disabled`: The application is stopped.
- Active
  - `N/A`: The agent is not enabled.
  - `Running`: The agent is enabled and responds to status requests.
  - `Degraded`: The agent is enabled and responds to status requests but its health is not `OK`.
  - `Failed`: The agent is enabled but does not respond to status requests.
- Version
- On-boot setting
  - `Enabled`: The application will start automatically at system startup (boot).
  - `Disabled`: The application will not start automatically.
- Details
- Application health
  - OK
  - Degraded
  - Fault
  - Unknown
  - N/A

8 Support and other resources

Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:

  https://www.hpe.com/info/assistance

- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:

  https://www.hpe.com/support/hpesc

Information to collect

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.
- To download product updates:

  Hewlett Packard Enterprise Support Center

  https://www.hpe.com/support/hpesc

  My HPE Software Center

  https://www.hpe.com/software/hpesoftwarecenter

- To subscribe to eNewsletters and alerts:

  https://www.hpe.com/support/e-updates

- To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:

  https://www.hpe.com/support/AccessToSupportMaterials

IMPORTANT Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Account set up with relevant entitlements.

Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which initiates a fast and accurate resolution based on the service level of your product. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

If your product includes additional remote support details, use search to locate that information.

HPE Get Connected

https://www.hpe.com/services/getconnected

HPE Tech Care Service

https://www.hpe.com/services/techcare

HPE Complete Care Service

https://www.hpe.com/services/completecare

Warranty information

To view the warranty information for your product, see the warranty check tool.

Regulatory information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the Hewlett Packard Enterprise Support Center:

https://www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional regulatory information

Hewlett Packard Enterprise is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements such as REACH (Regulation EC No 1907/2006 of the European Parliament and the Council). A chemical information report for this product can be found at:

https://www.hpe.com/info/reach

For Hewlett Packard Enterprise product environmental and safety information and compliance data, including RoHS and REACH, see:

https://www.hpe.com/info/ecodata

For Hewlett Packard Enterprise environmental information, including company programs, product recycling, and energy efficiency, see:

https://www.hpe.com/info/environment

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, use the **Feedback** button and icons (at the bottom of an opened document) on the Hewlett Packard Enterprise Support Center portal (https://www.hpe.com/support/hpesc) to send any errors, suggestions, or comments. This process captures all document information.

A Other management interfaces

SNMP reference

This topic describes the Simple Network Management Protocol (SNMP) capabilities that HPE MSA 2070/2072 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 HPE MSA 2070/2072 Storage systems is 1.3.6.1.4.1.11.2.51, where 51 is assigned for hpMSA. 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 `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 SMU.

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 is a subset of FA MIB 4.0, which is included with HPE System Insight Manager (SIM) and other products. The differences are described in .

FA MIB 2.2 was never formally adopted as a standard, but it is widely implemented and contains many elements useful for storage products. This MIB generally does not reference and integrate with other standard SNMP information. It is implemented under the experimental subtree.

Significant status within the device includes such elements as its temperature and power sensors, the health of its storage elements such as virtual disks, and the failure of any redundant component including an I/O controller. While sensors can be individually queried, for the benefit of network management systems all the above elements are combined into an "overall status" sensor. This is available as the unit status (`connUnitStatus` for the only unit).

The revisions of the various components within the device can be requested through SNMP.

The port section is only relevant to products with Fibre Channel (FC) host interface ports.

The event table allows 400 recently-generated events to be requested. Informational, minor, major, or critical event types can be selected. Whichever type is selected enables the capture of that type and more severe events. This mechanism is independent of the assignment of events to be generated into traps.

The traps section is not supported. It has been replaced by an ability to configure trap destinations using the CLI or the SMU. 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 10   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 10   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
connUnitConfiguration ChangeTime	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 SMU.	Default: info(8)
connUnitNumEvents	Number of events currently in the connUnitEventTable	Varies as the size of the Event Table varies

Table 10   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 page 113
`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 page 113
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 114
`connUnitSensorName`	Display string containing textual identification of the sensor intended primarily for operator use	See "External details for connUnitSensorTable" on page 114
`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 114
`connUnitSensor Characteristic`	Characteristics being monitored by this sensor	See "External details for connUnitSensorTable" on page 114
connUnitPortTable	**Includes the following objects as specified by the FA MIB2.2 Spec**	
`connUnitPortUnitId`	`connUnitId` of the connectivity unit that contains this port	Same as `connUnitId`
`connUnitPortIndex`	Unique value for each `connUnitPortEntry` between 1 and `connUnitNumPorts`	Unique value for each port, between 1 and the number of ports
`connUnitPortType`	Port type	not-present(3), or n-port(5) for point-to-point topology, or l-port(6)
`connUnitPortFCClassCap`	Bit mask that specifies the classes of service capability of this port. If this is not applicable, returns all bits set to zero.	Fibre Channel ports return 8 for class-three
`connUnitPortFCClassOp`	Bit mask that specifies the classes of service that are currently operational. If this is not applicable, returns all bits set to zero.	Fibre Channel ports return 8 for class-three

Table 10   FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
connUnitPortState	State of the port hardware	unknown(1), online(2), offline(3), bypassed(4)
connUnitPortStatus	Overall protocol status for the port	unknown(1), unused(2), ok(3), warning(4), failure (5), notparticipating(6), initializing(7), bypass(8)
connUnitPortTransmitter Type	Technology of the port transceiver	unknown(1) for Fibre Channel ports
connUnitPortModuleType	Module type of the port connector	unknown(1)
connUnitPortWwn	Fibre Channel World Wide Name (WWN) of the port if applicable	WWN octet for the port, or empty string if the port is not present
connUnitPortFCId	Assigned Fibre Channel ID of this port	Fibre Channel ID of the port All bits set to 1 if the Fibre Channel ID is not assigned or if the port is not present
connUnitPortSn	Serial number of the unit (for example, for a GBIC). If this is not applicable, returns an empty string.	Empty string
connUnitPortRevision	Port revision (for example, for a GBIC)	Empty string
connUnitPortVendor	Port vendor (for example, for a GBIC)	Empty string
connUnitPortSpeed	Speed of the port in KByte per second (1 KByte = 1000 Byte)	Port speed in KByte per second, or 0 if the port is not present
connUnitPortControl	Not supported	invalid(2) for an SNMP GET operation and not settable through an SNMP SET operation
connUnitPortName	String describing the addressed port	See
connUnitPortPhysical Number	Port number represented on the hardware	Port number represented on the hardware
connUnitPortStatObject	Not supported	0 (No statistics available)
connUnitEventTable	**Includes the following objects as specified by the FA MIB2.2 Spec**	
connUnitEventUnitId	connUnitId of the connectivity unit that contains this port	Same as connUnitId
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)
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

Table 10   FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
`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
Enterprise-specific fields	**Includes the following objects**	
`cpqSiSysSerialNum`	System serial number	For example, 3CL8Y40991
`cpqSiSysProductId`	System product ID	For example, 481321-001
`cpqSiProductName`	System product name	For example, HPE MSA 2070
`cpqHoMibStatusArray`	An array of MIB status structures. Octets 0-3 in block 0 are reserved for systems management and serve as an aggregate of the other MIBs.	Octet 0: 0. Octet 1 (overall status): 0 = Not available; 1 = Unknown/other; 2 = OK/normal; 3 = Degraded/warning; 4 = Failed/critical Octet 2 (system flags): 9 = device is not a server and web-based management is enabled Octet 3 (device type): 14 = enclosure For example, 00.02.09.14 (hex)
`cpqHoGUID`	Globally unique identifier formed from the product ID and serial number	For example, 4813213CL8Y40991

External details for certain FA MIB 2.2 objects

Tables in this topic specify values for certain objects described in Table 10.

External details for connUnitRevsTable

Table 11   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)
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 12 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)
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)

Table 12 connUnitSensorTable index, name, type, and characteristic values (continued)

connUnitSensorIndex	connUnitSensorName	connUnitSensorType	connUnitSensor Characteristic
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 13 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 SMU

1. Verify that the storage system's SNMP service is enabled. See "Enabling or disabling system-management services" on page 81.
2. Configure and enable SNMP traps. See "SNMP notifications" on page 90.
3. Optionally, configure a user account to receive SNMP traps. See "Managing SNMPv3 users" on page 85.

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 81. 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 83. The same users, security protocols, and passwords must be configured in the network management system.

Enterprise trap MIB

To download the HPE MSA Systems SNMP MIBs package, go to the download page at https://www.hpe.com/support/MSASNMPMIB.

Differences between FA 2.2 and 4.0 MIBs

FA MIB 2.2 is a subset of FA MIB 4.0. Therefore, SNMP elements implemented in MSA 2070/2072 systems can be accessed by a management application that uses FA MIB 4.0.

The following tables are *not* implemented in 2.2:

- `connUnitServiceScalars`
- `connUnitServiceTables`
- `connUnitZoneTable`
- `connUnitZoningAliasTable`
- `connUnitSnsTable`
- `connUnitPlatformTable`

The following variables are *not* implemented in 2.2:

- `connUnitFabricID`
- `connUnitNumLinks`
- `connUnitVendorId`
- `connUnitPortProtocolCap,`
  `connUnitPortProtocolOp,`
  `connUnitPortNodeWwn,`
  `connUnitPortHWState`
- `connUnitLinkCurrIndex`

Using SFTP/FTP

Although the SMU 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.

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 SFTP/FTP 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 SMU, prepare to use SFTP/FTP:

   a. Determine the network-port IP addresses of the system's controllers. See "Configuring controller network ports" on page 79.

   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 81.

   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 SFTP/FTP. See "User settings" on page 83.

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. 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`

6. 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.

7. 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 SFTP/FTP 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 40.

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 SMU, prepare to use SFTP/FTP:

   a. Determine the network-port IP addresses of the system's controllers. See "Configuring controller network ports" on page 79.

   b. Verify that the system's SFTP/FTP service is enabled. See "Enabling or disabling system-management services" on page 81.

   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 SFTP/FTP. See "User settings" on page 83.

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. 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.

6. 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 SMU, prepare to use SFTP/FTP:

   a. Determine the network-port IP addresses of the system's controllers. See "Configuring controller network ports" on page 79.

   b. Verify that the system's SFTP/FTP service is enabled. See "Enabling or disabling system-management services" on page 81.

   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 SFTP/FTP. See "User settings" on page 83.

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. 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`

6. In FTP, wait for the message `Operation Complete` to appear. No messages are displayed in SFTP; instead, the `get` command will return once the download is finished.

7. Quit the SFTP/FTP session.

Updating firmware

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

NOTE HPE recommends using the HPE Smart Component when updating firmware.

:bulb: **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. Storage systems in a replication set should run the same or compatible firmware versions. You can update the firmware in each controller module by loading a firmware file obtained from the HPE web download site, https://www.hpe.com/storage/msafirmware. To install a firmware Smart Component, follow the instructions on the HPE website. Otherwise, to install a firmware binary file, follow the steps below.

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 SMU, 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 SFTP/FTP.
3. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the directory containing the firmware file to load.

4. Enter:

sftp -P \<port\> \<controller-network-address\>

or

ftp \<controller-network-address\>

For example:

sftp -P 1022 10.235.216.152

ftp 10.1.0.9

5. Log in as an SFTP/FTP user.

6. 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. If the controller enclosure has attached enclosures, allow additional time for each expansion module's enclosure management processor (EMP) to be updated. This typically takes 2.5 minutes for each EMP in an MSA 2070/2072 disk enclosure.

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 SMU 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 SMU 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.

7. Quit the SFTP/FTP session.

8. Clear your web browser's cache, then sign in to the SMU. 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 expansion-module firmware

Expansion-module firmware is included with the main controller enclosure firmware, and cannot be installed separately for MSA 2070/2072 Storage systems.

Updating disk firmware

You can update disk firmware by loading a firmware file obtained from the HPE web download site, https://www.hpe.com/storage/msafirmware. To install a firmware Smart Component, follow the instructions on the HPE website.

Installing a license file

1. In the SMU, 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 SFTP/FTP.
2. Ensure that the license file is saved to a network location that the storage system can access.
3. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the directory containing the license file to load.
4. Log in to the controller enclosure that the file was generated for:

   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
   ```
5. Log in as an SFTP/FTP user.
6. Enter:

   put <license-file> **license**

   For example:

   ```
   put certificate.txt license
   ```
   In FTP, a message confirms whether installation succeeded or failed. If installation succeeds, licensing changes take effect immediately. No messages are displayed in SFTP.

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. 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 user certificate and its associated trust certificates per controller. Upload the root trust certificate first, then any intermediate trust certificates, and finally the device certificate. This allows the trust certificate that issued them to verify intermediate trust or device certificates. To upload multiple user certificates on a controller, use the `add certificate` API command as described in the CLI Reference Guide.

Trust certificates are maintained separately by each controller. If you load a trust certificate to one controller, it is not automatically copied to the partner controller. You must upload both trust and user certificates separately to each controller.

NOTE   Details pertaining to certificates:

- By default, the system generates a unique SSL certificate for each controller. When using SFTP/FTP to install certificates, you must explicitly install the certificate to the controller for which the session applies. For example, you cannot install a certificate to controller B during an SFTP/FTP session with controller A. You cannot install a certificate to both controllers from a given session: the installation tasks are discrete.
- The supported file format for certificates is PEM (Privacy-Enhanced Mail) encoded only.
- Supported certificate versions are x.509 v1 and v3. Avoid using unsupported versions of certificates.
- Extensions are supported if uploading certificates through SFTP/FTP.
- The key generated by running the `create certificate-signing-request` CLI command has a length of 2048 bits if the encryption type is RSA or 256 bits if the encryption type is ECC.

To install a security certificate

1. In the SMU, prepare to use SFTP/FTP:
   a. Determine the network-port IP addresses of the system's controllers. See "Configuring controller network ports" on page 79.
   b. Verify that the system's SFTP/FTP service is enabled. See "Enabling or disabling system-management services" on page 81.
   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 SFTP/FTP. See "Managing local users" on page 83.
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. Enter:

   put \<certificate-filename\> **cert-file:**{**usr**|**trust**}

where:

- <certificate-filename>: Name of the certificate file for your specific system.
- `usr`: Device/server certificate uploaded by end user.
- `trust`: Trusted intermediate or root certificate.

The certificate is automatically activated.

6. In the CLI:
   a. Run `restart mc full` to restart the Management Controller and use the uploaded certificate.
   b. Run `show certificates` to view the new certificate.
7. Because restarting the Management Controller restarted the internal HTTPS server, log out of any sessions and log in again for the new security certificate to take effect.

Downloading system heat map data

If requested by support engineers for analysis, you can download cumulative I/O density data, also known as heat map data, from the system.

To gather this data, access the storage system's SFTP/FTP interface and use the `get logs` command with the `heatmap` option to download a log file in CSV format. The file contains data for the past 7 days from both controllers.

1. In the SMU, prepare to use SFTP/FTP:
   a. Determine the network-port IP addresses of the system's controllers. See "Configuring controller network ports" on page 79.
   b. Verify that the system's SFTP/FTP service is enabled. See "Enabling or disabling system-management services" on page 81.
   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 SFTP/FTP. See "Managing local users" on page 83.
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 that has permission to use the SFTP/FTP interface.
5. Enter:

   get logs:heatmap <filename>**.csv**

   where: <filename> is the file that will contain the data.

   For example:

   ```
   get logs:heatmap IO_density.csv
   ```

   In FTP, wait for the message `Operation Complete` to appear. No messages are displayed in SFTP; instead, the `get` command will return once the download is finished.
6. Quit the SFTP/FTP session.

Using SLP

HPE MSA 2070/2072 Storage systems support Service Location Protocol (SLP, srvloc), which is a service discovery protocol that allows computers and other devices to find services in a LAN without prior configuration. SLP is open for use on all operating systems, and does not require formal licensing.

SLP is based on User Datagram Protocol (UDP) and can use Transmission Control Protocol (TCP) if needed. SLP listens on port 427. When a client, or User Agent (UA), connects to a network, the client queries for Directory Agents (DA) on the network. If no DA responds, the client assumes a DA-less network and sends a multicast UDP query. All Service Agents (SA) that contain query matches will send a UDP answer to the client. If the answer message is too large, the client can repeat the query using TCP.

In a network with DAs, each SA must register all services with a DA. Then the clients will query the DAs, who will respond to the query with its cached SA information.

Through use of DAs, SLP can also scale beyond the local area network to large enterprise, which is an enterprise IT issue. Consult the IETF RFC2165.

When SLP is enabled, the storage system will advertise the interfaces shown in Table 14 and populate the configuration attributes shown in Table 15.

Table 14   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`

SLP attributes are listed below.

Table 15   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

Table 15 SLP attributes shown for a storage system (continued)

SLP attribute	Corresponding property shown by the CLI `show system detail` command in XML API mode
`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 SMU, as described in "Enabling or disabling system-management services" on page 81, 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.

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 40.

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 117.

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 16   System information defaults

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

Table 17   Management protocols defaults

Setting	Value
CLI/Telnet	Disabled
CLI/SSH	Enabled
SLP	Disabled
FTP	Disabled
SFTP	Enabled
SNMP	Disabled
WBI/HTTP	Disabled
WBI/HTTPS	Enabled
Debug	Disabled
Ciphers setting	Default cipher strings

Table 18   User defaults

Setting	Value
Users	All configured users are deleted and replaced with default user definitions and default settings: User: `setup`; Password: press ENTER

Table 19   SNMP defaults

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

Table 20   SMTP defaults

Setting	Value
Email notification	Disabled
Email notify filter	None
Email addresses	None
Email server	None

Table 20 SMTP defaults (continued)

Setting	Value
Email domain	None
Email sender	None
Log destination	None
Include logs	Disabled
Persistent alerts	Enabled
Event notification	None
Alert notification	All
Proxy setting	Cleared

Table 21 LDAP defaults

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

Table 22 Syslog defaults

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

Table 23 Host port defaults

Setting	Value
FC link speed	Auto
FC topology	Point-to-point

Table 24 Disk spin down defaults

Setting	Value
Disk spin down	Disabled

Table 25 Advanced defaults

Setting	Value
Disk group background scrub	Enabled
Partner firmware upgrade	Enabled
Utility priority	High
SMART	Enabled
Dynamic spare configuration	Enabled
Enclosure polling rate	5 seconds
Host control of caching	Disabled
Sync cache mode	Immediate

Table 25 Advanced defaults (continued)

Setting	Value
Missing LUN response	Illegal Request
Controller failure	Enabled
Supercap failure	Enabled
Power supply failure	Disabled
Fan failure	Disabled
Temperature exceeded	Disabled
Partner notify	Enabled
Auto write back	Enabled
Inactive drive spin down	Disabled
Inactive drive spin down delay	15 minutes
Managed logs	Disabled
Auto stall recovery	Enabled (for failover/failback, not I/O)
Restart on CAPI fail	Enabled
Default mapping	Disabled
Slow disk detection	Enabled
Scrub schedule	Disabled

Table 26 FDE defaults

Setting	Value
FDE settings	Preserved

Table 27 Replication defaults

Setting	Value
Peer connections	Preserved
Replication sets	Preserved
CHAP records	Preserved

Table 28 Enclosure defaults

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

Table 29 iSCSI port defaults

Setting	Value
IP	Preserved
IP version	Preserved
Netmask	Preserved
Gateway	Preserved
Router (IPv6 only)	Preserved

Table 30 Other iSCSI defaults

Setting	Value
CHAP enabled	Preserved
iSNS	Preserved
Jumbo frames	Preserved

Table 31 Host defaults

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

Table 32 Volume defaults

Setting	Value
Volume identifying information	Preserved
Volume groups	Preserved

Table 33 Pool defaults

Setting	Value
Thresholds	Preserved
Overcommit	Preserved
Limits and policy	Preserved
Snapshot space thresholds	Preserved

Table 34 Other defaults

Setting	Value
CLI parameters	CLI parameters are kept on a per-user basis. All configured users are deleted and replaced with default user definitions and default settings as detailed in the Users section of this table.
CLI session timeout	Preserved
Tasks and schedules	Preserved
Debug log parameters	Each parameter is reset to its default as documented for the `set debug-log-parameters` CLI command.
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
Alert condition history	Preserved

Table 34 Other defaults (continued)

Setting	Value
Alerts	Preserved
SSL/SSH certificates	Preserved
Licenses	Preserved
Disk group metadata	Preserved
Volume snapshot retention priority	Preserved
Volume cache settings	Preserved
Expander PHY settings	Controller module root expander PHY settings are cleared
Volume tier affinity	Preserved
Device identification LED status	Cleared

D System configuration limits

Table 35 System configuration limits

Feature	Value
Enclosures and disks	
Maximum enclosures and disks per system	Supported configurations:
2070/2072	One 2U12 controller enclosure + nine 2U12 expansion enclosures = 120
	One 2U12 controller enclosure + nine 2U24 expansion enclosures = 228
	One 2U24 controller enclosure + nine 2U12 expansion enclosures = 132
	One 2U24 controller enclosure + nine 2U24 expansion enclosures = 240
Disk groups and pools	
Storage model	Virtual
Maximum pools per controller module	1
Maximum usable pool size	4096TiB (4PiB)
Maximum non-MSA-DP+ disk-group size	4096TiB (4PiB)
Maximum MSA-DP+ disk-group size	1536TiB (1.5PiB)
Maximum disk groups per pool	16
Maximum disk groups per controller module	16
Minimum/maximum disks per disk group	NRAID (non-RAID): 1/1 (read cache only) RAID 0: 2/2 (read cache only) RAID 1: 2/2 RAID 5: 3/16 RAID 6: 4/16 RAID 10: 4/16 MSA-DP+: 12/128
Maximum global spares per system	64
Maximum MSA-DP+ disk groups per controller module	4
MSA-DP+ stripe width (data+parity)	8+2
Volumes, initiators, hosts, and mapping	
Maximum volumes per system	1024
Maximum volume (LUN) size	128TiB (140TB)
Maximum mappable volumes (LUNs) per pool	512
Maximum volumes per pool	1024 (512 base volumes and 512 snapshots)
Maximum volumes per volume group	1024
Maximum volume groups per system	256
Maximum volumes per replication volume group	16
Maximum hosts per system	512
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

Table 35 System configuration limits (continued)

Feature	Value
Maximum host groups per system	32
Maximum queue depth per host port	1024
Maximum FC host-port link speed	16Gb/s
Maximum iSCSI host-port link speed	25Gb/s
Maximum SAS host-port link speed	12Gb/s
Virtual volume snapshots	
Maximum snapshots per system, without a license	64
Maximum snapshots per system, with a license	512
Maximum snapshots per pool (net usable)	512
Maximum base volumes per system	1024
Maximum snapshots per base volume	254 in the volume's snapshot tree (with a license)
Maximum mappable snapshots per system	512
Virtual volume replication	
Maximum peer connections per system	4
Maximum replicated volumes per system	32
Maximum replication sets per volume	1
Maximum volumes for a replicated volume group	16, if no other volumes belong to a replication set
Minimum replication schedule interval	30 minutes
Miscellaneous	
Maximum SCSI reservations per system	1024
Maximum SCSI reservations per LUN	1
Maximum SCSI registrations per system	32768
Maximum SCSI registrations per LUN	4096
Maximum management sessions per controller	20

Glossary

2U12

An enclosure that is two rack units in height and can contain 12 disks.

2U24

An enclosure that is two rack units in height and can contain 24 disks.

AES

Advanced Encryption Standard.

AFA

All-flash array (AFA). A storage system that uses only SSDs, without tiering.

all-flash array

All-flash array (AFA). A storage system that uses only SSDs, without tiering.

allocated page

A page of virtual pool space that has been allocated to a volume to store data.

allocation rate

The rate, in pages per minute, at which a virtual pool is allocating pages to its volumes because they need more space to store data.

array

Synonym for storage system.

ASC/ASCQ

Additional Sense Code/Additional Sense Code Qualifier. Information on sense data returned by a SCSI device.

auto-write-through

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

automated tiered storage

Automated tiered storage. A virtual-storage feature that automatically uses the appropriate tier of disks to store data based on how frequently the data is accessed. This enables higher-cost, higher-speed disks to be used only for frequently needed data, while infrequently needed data can reside in lower-cost, lower-speed disks.

available disk

A disk that is not a member of a disk group, is not configured as a spare, and is not in the leftover state. It is available to be configured as a part of a disk group or as a spare.

AWT

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

base volume

A virtual volume that is not a snapshot of any other volume, and is the root of a snapshot tree.

base volume allocated

The amount of thin-provisioned virtual storage used for volume data. This is represented in GB/TB and as a percentage of the total physical storage capacity.

canister

Synonym for IOM.

CAPI

Configuration Application Programming Interface. A proprietary protocol used for communication between the Storage Controller and the Management Controller in a controller module. CAPI is always enabled.

CHAP

Challenge Handshake Authentication Protocol.

chassis

The sheet metal housing of an enclosure.

child volume

The snapshot of a parent volume in a snapshot tree.

chunk size

The amount of contiguous data that is written to a disk group member before moving to the next member of the disk group.

compatible disk

A disk that can be used to replace a failed member disk of a disk group because it has at least the same capacity as, and is of the same type (enterprise SAS, for example) as, the disk that failed.

controller A (or B)

A short way of referring to controller module A (or B).

controller enclosure

An enclosure that 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.

deallocation rate

The rate, in pages per minute, at which a pool is deallocating pages from its volumes because they no longer need the space to store data.

default mapping

Access settings for a volume that apply to all current and future initiators and use a single LUN.

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.

disk spin down

Disk spin down (DSD). A power-saving feature for spinning disks that monitors disk activity and spins down inactive disks based on user-selectable policies.

DNS

Domain Name System.

drain

The automatic movement of active volume data from a virtual disk group to other disk-group members within the same pool.

drive enclosure

Synonym for expansion enclosure. See also EBOD, JBOD.

drive spin down

Disk spin down (DSD). A power-saving feature for spinning disks that monitors disk activity and spins down inactive disks based on user-selectable policies.

DSD

Disk spin down (DSD). A power-saving feature for spinning disks that monitors disk activity and spins down inactive disks based on user-selectable policies.

DSP

Digital signal processor.

dual-port disk

A disk that is connected to both controllers so it has two data paths, achieving fault tolerance.

dynamic spare

An available compatible disk that is automatically assigned, if the dynamic spares option is enabled, to replace a failed disk in a disk group with a fault-tolerant RAID level.

EBOD

Expanded Bunch of Disks. Expansion enclosure attached to a controller enclosure.

EC

Expander Controller (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.

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.

expansion enclosure

An enclosure that can contain two expansion modules. Expansion enclosures can be connected to a controller enclosure to provide additional storage capacity. See also EBOD, JBOD.

expansion module

A FRU that contains the following subsystems and devices: a SAS expander and Expander Controller processor; host, expansion, and service ports; and midplane connectivity. In an expansion enclosure, the upper expansion module is designated A and the lower one is designated B.

explicit mapping

Access settings for a volume to a specific initiator, host, or host group using a unique LUN.

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.

FC

Fibre Channel interface protocol.

FC-AL

Fibre Channel Arbitrated Loop. The FC topology in which devices are connected in a one-way loop.

FDE

Full Disk Encryption (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.

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 can be either a controller module or an expansion module.

initiator

An external port to which the storage system is connected. The external port may be a port in an I/O adapter in a server, or a port in a network switch.

IOM

Input/output module (I/O module, IOM). An IOM can be either a controller module or an expansion module.

IOPS

I/O operations per second.

IQN

iSCSI Qualified Name.

iSCSI

Internet SCSI interface protocol.

iSNS

Internet Storage Name Service.

JBOD

"Just a bunch of disks." An expansion enclosure attached to a server.

LBA

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

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.

LIP

Loop Initialization Primitive. An FC primitive used to determine the loop ID for a controller.

lock key

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

loop

Fibre Channel Arbitrated Loop. The FC topology in which devices are connected in a one-way loop.

LUN

Logical Unit Number. A number that identifies a mapped volume to a host system.

MAC address

Media Access Control Address. A unique identifier assigned to network interfaces for communication on a network.

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.

MAS

MC Auxiliary Services.

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 and snapshots in the disk group, host mapping of volumes, and results of the last media scrub.

MIB

Management Information Base. A database used for managing the entities in SNMP.

midplane

The printed circuit board to which components connect in the middle of an enclosure.

mount

To enable access to a volume from a host OS. Synonyms for this action include present and map.

MSA-DP+

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.

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.

OID

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

orphan data

See unwritable cache data.

overcommit

A setting that controls whether a pool is allowed to have volumes whose total size exceeds the physical capacity of the pool.

overcommitted

The amount of storage capacity that is allocated to virtual volumes exceeds the physical capacity of the storage system.

page

A range of contiguous LBAs in a virtual disk group.

paged storage

A method of mapping logical host requests to physical storage that maps the requests to virtualized "pages" of storage that are in turn mapped to physical storage. This provides more flexibility for expanding capacity and automatically moving data than the traditional, linear method in which requests are directly mapped to storage devices. Paged storage is also called virtual storage.

parent volume

A virtual volume that has snapshots (can be either a base volume or a base snapshot volume). The parent of a snapshot is its immediate ancestor in the snapshot tree.

partner firmware update

Partner firmware update (PFU). The automatic update of the partner controller when the user updates firmware on one controller.

passphrase

A user-created password that allows users to manage lock keys in an FDE-capable system. Contains up to 32 characters and stored in non-volatile memory. See also FDE, lock key.

PCB

Printed circuit board.

PCBA

Printed circuit board assembly.

PCM

Power and cooling module FRU. A power supply module that includes an integrated fan. See also PSU.

PDU

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

peer connection

The configurable entity defining a peer-to-peer relationship between two systems for the purpose of establishing an asynchronous replication relationship. See also peer system.

peer system

A remote storage system that can be accessed by the local system and is a candidate for asynchronous replications. Both systems in a peer connection are considered peer systems to each other, and they both maintain a peer connection with the other. Asynchronous replication of volumes may occur in either direction between peer systems configured in a peer connection.

PFU

Partner firmware update (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.

point-to-point

Fibre Channel Point-to-Point topology in which two ports are directly connected.

pool

A container for volumes that is composed of one or more virtual disk groups.

POST

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

primary system

The storage system that contains a replication set's primary volume.

primary volume

The volume that is the source of data in a replication set and that can be mapped to hosts. The primary volume exists in the primary storage system.

provisioning

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

RAID head

Synonym for controller enclosure.

RBOD

"RAID bunch of disks." See controller enclosure.

read cache

A special virtual disk group, comprised of SSDs, that can be added to a pool for the purpose of speeding up read access to data stored on spinning disks elsewhere in the pool.

recovery

In an active-active configuration, recovery is the act of returning ownership of controller resources to a controller (which was offline) from its partner controller. The resources include volumes, cache data, host ID information, and LUNs and WWNs. See also failover.

replication

Asynchronous replication of block-level data from a volume in a primary system to a volume in a secondary system by creating an internal snapshot of the primary volume and copying the snapshot data to the secondary system via FC or iSCSI links.

replication set

A container that houses the internal support for performing replication. It defines a relationship between a primary and secondary volume for the purposes of maintaining a remote copy of the primary volume on a peer system. See primary volume, secondary volume.

replication set failover

The replication set's secondary system has allowed direct access to the secondary volume or volume group because the primary system is not operational. In this state no replications will occur, even if the primary system becomes operational and communication is restored. The secondary volume can be mapped and accessed for use, including rollback to the contents of any manually created or snapshot-history snapshot.

replication snapshot history

As part of handling a replication, the replication set will automatically take a snapshot of the primary and/or secondary volume, thereby creating a history of data that has been replicated over time. This feature can be enabled for a secondary volume or for a primary volume and its secondary volume.

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.

SAS

Serial Attached SCSI (Small Computer System Interface).

SATA

Serial Advanced Technology Attachment.

SC

Storage Controller. A processor (located in a controller module) that is responsible for RAID controller functions. The SC is also referred to as the RAID controller. See also EC, MC.

secondary system

The storage system that contains a replication set's secondary volume. See also primary system.

secondary volume

The volume that is the destination for data in a replication set and that is not accessible to hosts. The secondary volume exists in a secondary storage system.

secret

For use with CHAP, a password that is shared between an initiator and a target to enable authentication.

SED

Self-encrypting drive. A disk drive that provides hardware-based data encryption and supports 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. 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.

SMU

Storage Management Utility 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.

snapshot

A point-in-time copy of the data in a source volume that preserves the state of the data as it existed when the snapshot was created. Data associated with a snapshot is recorded in the source volume. A snapshot can be mapped and written to. Snapshots that can be mapped to hosts are counted against the snapshot-license limit, whereas transient and unmappable snapshots are not.

snapshot allocated

The amount of thin-provisioned virtual storage allocated for all snapshot data. This is represented in GB/TB and as a percentage of the total physical storage capacity.

snapshot tree

A group of virtual volumes that are interrelated due to creation of snapshots. Since snapshots can be taken of existing snapshots, volume inter-relationships can be thought of as a "tree" of volumes. A tree can be 254 levels deep. See also base volume, child volume, parent volume, source volume.

source volume

A volume that has snapshots. Used as a synonym for parent volume.

SSD

Solid-state drive.

SSH

Secure Shell. A network protocol for secure data communication.

SSL

Secure Sockets Layer. A cryptographic protocol that provides security over the internet.

standard volume

A volume that can be mapped to initiators and presented as a storage device to a host system, but is not enabled for snapshots.

Storage Controller

Storage Controller. A processor (located in a controller module) that is responsible for RAID controller functions. The SC is also referred to as the RAID controller. See also EC, MC.

Storage Management Utility

Storage Management Utility 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, optionally with connected expansion enclosures. Product documentation and interfaces use the terms storage system and system interchangeably.

syslog

A protocol for sending event messages across an IP network to a logging server. This feature supports User Datagram Protocol (UDP) but not Transmission Control Protocol (TCP).

TCP

Transmission Control Protocol.

thin provisioning

A virtual-storage feature that allows actual storage for a volume to be assigned as data is written, rather than storage being assigned immediately for the eventual size of the volume. This allows the storage administrator to overcommit physical storage, which in turn allows the connected host system to operate as though it has more physical storage available than is actually allocated to it. When physical resources fill up, the storage administrator can add storage capacity on demand.

tier

A homogeneous group of disks, typically of the same capacity and performance level, that comprise one or more virtual disk groups in the same pool. Tiers differ in their performance, capacity, and cost characteristics, which forms the basis for the choices that are made with respect to which data is placed in which tier. The predefined tiers are: Performance, which uses SSDs; Standard, which uses enterprise-class spinning SAS disks; Archive, which uses midline spinning SAS disks.

tier migration

The automatic movement of blocks of data, associated with a single virtual volume, between tiers based on the access patterns that are detected for the data on that volume.

total committed

The total storage capacity assigned to volumes. For example, 10 volumes with a capacity of 1TB will result in 10TB of committed storage space.

tray

Synonym for enclosure.

UDP

User Datagram Protocol.

ULP

Unified LUN Presentation. A RAID controller feature that enables a host system to access mapped volumes through any controller host port. ULP incorporates ALUA extensions.

undercommitted

The amount of storage capacity that is allocated to volumes is less than the physical capacity of the storage system.

unmount

To remove access to a volume from a host OS. Synonyms include unpresent and unmap.

unused capacity

The remaining storage capacity in the pool for creating new volumes or snapshots. This is represented as the pool capacity minus the committed capacity.

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 SMU and CLI interfaces.

virtual

The storage-class designation for logical components such as volumes that use paged-storage technology to virtualize data storage. See paged storage.

volume

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

volume copy

An independent copy (clone) of the data in a virtual volume. The capability to copy volumes makes use of snapshot functionality.

volume group

A user-defined group of volumes for ease of management, such as for host-attachment operations.

VPD

Vital Product Data. Data held on an EEPROM in an enclosure or FRU that is used by GEM to identify and control the component.

WBI

Storage Management Utility 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

S