



Hewlett Packard
Enterprise

HPE MSA 1050/2050 SMU Reference Guide

Abstract

This guide is for use by storage administrators to manage an HPE MSA 1050/2050 storage system by using its web interface, Storage Management Utility (SMU).

Firmware Version: V270

Part Number: Q1J79-62025

Published: July 2018

Edition: 1

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Intel®, Itanium®, Pentium®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the United States and other countries.

Microsoft® and Windows® are U.S. trademarks of the Microsoft group of companies.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java and Oracle are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

Revision History

822372-002	December 2015
------------	---------------

Initial HPE release

Q1J79-62008	September 2017
-------------	----------------

Updated for MSA 2050 SAS, 1050 SAN, and 1050 SAS.

Q1J79-62025	July 2018
-------------	-----------

Updated for MSA 2050 SAS, 1050 SAN, and 1050 SAS.

Contents

1 Getting started	10
Configuring and provisioning a new storage system	10
Using the interface	11
Web browser requirements and setup	11
Areas of the interface	11
Tips for using the SMU	12
Tips for using tables	12
Exporting data to a CSV file	13
Tips for using help	13
Color codes	14
Size representations	15
Signing in and signing out	15
System concepts	16
About virtual storage	16
About the handling of linear storage	16
About disk groups	17
About RAID levels	19
About SSDs	20
About SSD read cache	21
About spares	22
About pools	22
About volumes and volume groups	23
About volume cache options	23
About thin provisioning	24
About automated tiered storage	25
About initiators, hosts, and host groups	25
About volume mapping	26
About snapshots	27
About copying volumes	28
About reconstruction	28
About quick rebuild	29
About performance statistics	29
About firmware update	30
About managed logs	30
About LDAP	31
About replicating virtual volumes	34
About the Full Disk Encryption feature (for MSA 2050 only)	35
About data protection with a single controller	35
About SAS cabling (for MSA 1050 only)	36
2 Working in the Home topic	37
Using guided setup	37
Viewing overall system status	37
Host information	37
Port information	38
Capacity information	39
Storage information	39
System health information	40

Spares information	40
Resolving a pool conflict caused by inserting a foreign disk group	40
Configuring system settings	41
Setting the system date and time	41
Managing users	42
Installing a license	48
Configuring controller network ports	49
Enabling or disabling system-management services	50
Entering system identification information	52
Setting system notification settings	52
Changing host port settings	56
Managing scheduled tasks	59
3 Working in the System topic	61
Viewing system components	61
Front view	61
Rear view	61
Table view	62
Configuring system settings	64
Resetting host ports	64
Rescanning disk channels	64
Clearing disk metadata	65
Updating firmware	66
Best practices for firmware update	66
Updating controller module firmware	66
Updating expansion module firmware	67
Updating disk-drive firmware	68
Using the activity progress interface	69
Changing FDE settings (for MSA 2050 only)	70
Changing FDE general configuration	70
Repurposing the system	72
Repurposing disks	72
Setting import lock key IDs	72
Configuring advanced settings	73
Changing disk settings	73
Changing system cache settings	75
Configuring partner firmware update	76
Configuring system utilities	77
Restarting or shutting down controllers	78
Restarting controllers	78
Shutting down controllers	79
4 Working in the Hosts topic	80
Viewing hosts	80
Hosts table	80
Related Maps table	80
Creating an initiator	81
Modifying an initiator	81
Deleting initiators	81
Adding initiators to a host	82
Removing initiators from hosts	82

Removing hosts	82
Renaming a host	82
Adding hosts to a host group	83
Removing hosts from a host group	83
Renaming a host group	83
Removing host groups	83
Configuring CHAP	84
5 Working in the Pools topic	86
Viewing pools	86
Pools table	86
Related Disk Groups table	86
Related Disks table	87
Adding a disk group	88
Add Disk Group panel overview	88
Virtual disk groups	89
Read-cache disk groups	89
Disk group options	90
Modifying a disk group	90
Removing disk groups	91
Managing Spares	92
Global spares	92
Creating a volume	93
Changing pool settings	93
Verifying and scrubbing disk groups	94
Verifying a disk group	94
Scrubbing a disk group	95
6 Working in the Volumes topic	96
Viewing volumes	96
Volumes table	96
Snapshots table	96
Maps table	97
Replication Sets table	97
Schedules table	98
Creating a virtual volume	99
Modifying a volume	100
Copying a volume or snapshot	100
Aborting a volume copy	101
Adding volumes to a volume group	101
Removing volumes from a volume group	101
Renaming a volume group	102
Removing volume groups	102
Rolling back a volume	103
Deleting volumes and snapshots	103
Creating snapshots	104
Resetting a snapshot	105

Creating a replication set from the Volumes topic	105
Primary volumes and volume groups	106
Secondary volumes and volume groups	106
Queuing replications.	106
Maintaining replication snapshot history from the Volumes topic.....	106
Initiating or scheduling a replication from the Volumes topic.....	108
Managing replication schedules from the Volumes topic.....	110
7 Working in the Mappings topic.....	111
Viewing mappings	111
Mapping initiators and volumes	111
Removing mappings.....	114
Removing all mappings.....	114
Viewing map details	114
8 Working in the Replications topic	116
About replicating virtual volumes.....	116
Replication prerequisites	116
Replication process.....	117
Creating a virtual pool for replication.....	119
Setting up snapshot space management in the context of replication	119
Replication and empty allocated pages.....	119
Disaster recovery.....	120
Replication licensing	121
Viewing replications	121
Peer Connections table	121
Replication Sets table.....	122
Replication Snapshot History table.....	122
Querying a peer connection	123
Creating a peer connection.....	123
CHAP and replication.....	124
Modifying a peer connection	125
Deleting a peer connection	126
Creating a replication set from the Replications topic.....	126
Primary volumes and volume groups	126
Secondary volumes and volume groups	127
Queuing replications.....	127
Maintaining replication snapshot history from the Replications topic.....	127
Modifying a replication set	129
Deleting a replication set.....	130
Initiating or scheduling a replication from the Replications topic	130
Aborting a replication.....	132
Suspending a replication	132
Resuming a replication.....	133
Managing replication schedules from the Replications topic.....	133
9 Working in the Performance topic.....	135
Viewing performance statistics	135
Historical performance graphs	135
Updating historical statistics	137

Exporting historical performance statistics	137
Resetting performance statistics	138
10 Working in the banner and footer.....	139
Banner and footer overview.....	139
Viewing system information.....	139
Viewing certificate information.....	139
Viewing connection information.....	140
Viewing system date and time information	140
Changing date and time settings.....	140
Viewing user information	141
Viewing health information	141
Saving log data to a file	141
Viewing event information	142
Viewing the event log	143
Viewing capacity information	143
Viewing host I/O information.....	144
Viewing tier I/O information.....	144
Viewing I/O workload activity	144
Reading the graph	145
Viewing recent system activity	145
Viewing the notification history.....	145
11 Support and other resources	147
Accessing Hewlett Packard Enterprise Support	147
Information to collect	147
Accessing updates.....	147
Customer self repair	148
Remote support	148
Remote support and Proactive Care information.....	148
Proactive Care customer information	148
Warranty information.....	148
Additional warranty information	148
Regulatory information	149
Additional regulatory information.....	149
Documentation feedback	149
A Other management interfaces.....	150
SNMP reference.....	150
Supported SNMP versions.....	150
Standard MIB-II behavior	150
Enterprise traps.....	150
FA MIB 2.2 SNMP behavior	151
External details for certain FA MIB 2.2 objects.....	156
Configuring SNMP event notification in the SMU	159
SNMP management	159
Enterprise trap MIB	160
FA MIB 2.2 and 4.0 differences.....	162
Using FTP and SFTP.....	163
Downloading system logs	163
Transferring log data to a log-collection system	164

Downloading historical disk-performance statistics	165
Updating firmware	166
Installing a license file	171
Installing a security certificate	171
Downloading system heat map data	172
Using SMI-S	172
Embedded SMI-S array provider	173
About the MSA 1050/2050 SMI-S provider	174
SMI-S profiles	175
CIM	176
Life cycle indications	177
SMI-S configuration	178
Listening for managed-logs notifications	179
Testing SMI-S	179
Troubleshooting	180
Using SLP	180
B Administering a log-collection system	182
How log files are transferred and identified	182
Log-file details	182
Storing log files	182
Glossary	184
Index	192

Tables

1	Areas of the SMU interface.....	11
2	Home topic storage space color codes.....	14
3	Create Virtual Volumes panel storage space color codes.....	15
4	Storage size representations in base 2 and base 10.....	15
5	Decimal (radix) point character by locale	15
6	Example applications and RAID levels	19
7	RAID level comparison.....	19
8	Number of disks per RAID level to optimize virtual disk group performance	20
9	Settings for the default users.....	43
10	Additional information for rear view of enclosure.....	62
11	Activity progress properties and values.....	69
12	Available host groups, hosts, and initiators.....	112
13	Available volume groups and volumes.....	113
14	Historical performance graphs	135
15	Connection information.....	140
16	FA MIB 2.2 objects, descriptions, and values.....	151
17	connUnitRevsTable index and description values	156
18	connUnitSensorTable index, name, type, and characteristic values	158
19	connUnitPortTable index and name values	159
20	Supported SMI-S profiles.....	175
21	CIM Alert indication events.....	177
22	Life cycle indications.....	177
23	CLI commands for SMI-S protocol configuration.....	179
24	Troubleshooting.....	180
25	Interfaces advertised by SLP	181
26	SLP attributes shown for a storage system	181

1 Getting started

The Storage Management Utility (SMU) is a web-based application for configuring, monitoring, and managing the storage system. The SMU is a web-based 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. 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 SNMP, FTP, SFTP, SMI-S, SLP, and command-line (CLI) interfaces. For information about all interfaces other than the CLI, see this guide. For information about using the CLI, see the CLI Reference Guide.

Product information specific to HPE MSA 1050 or 2050 is called out by the product name in bold at the beginning of each relevant paragraph (for instance, “For MSA 2050:”). If there is a mention within a sentence, there will be a callout within the sentence. If there are several paragraphs for a product, a preceding note will indicate the relevant content. If an entire section is specific to a product, the section heading will include a call out, which will apply to all sub-sections. Otherwise, the content of this guide applies to both systems.

Configuring and provisioning a new storage system

The SMU provides options for you to quickly and easily set up your system by guiding you through the firmware update and configuration process. When configuring your system, you must first access the Update Firmware panel where you can review the controller module firmware version and perform recommended updates. When finished, you must configure your system settings by accessing the System Settings panel and completing all required options. Once these options are complete you can provision your system.

To setup and configure your system

1. Configure your web browser to use the SMU as described in [“Web browser requirements and setup” \(page 11\)](#).
2. Sign in to the SMU; the default user for management is `manage` and the default password is `!manage`. The Welcome panel displays.
For more information about signing in, see [“Signing in and signing out” \(page 15\)](#). For more information about the Welcome panel, see [“Using guided setup” \(page 37\)](#).
3. From the Welcome panel, click **Upgrade Firmware**.
4. Review the controller module firmware version.
5. Per HPE recommendation, verify your firmware is up to date as described in [“Updating firmware” \(page 66\)](#).
6. When you are finished with the Update Firmware panel, click **Close**.
7. From the Welcome panel, click **System Settings**.
8. Choose options to configure your system. For more information, see [“Configuring system settings” \(page 41\)](#).

NOTE: Tabs with a red asterisk next to them are required.

9. Save your settings and exit System Settings to return to the Welcome panel.
10. Create disk groups and pools, as described in [“Adding a disk group” \(page 88\)](#).
11. Create volumes and volume groups and map them to hosts, host groups, and initiators, as described in [“Creating a volume” \(page 93\)](#).
12. From hosts, verify volume mappings by mounting the volumes and performing read/write tests to the volumes.
13. Optionally, for replication of virtual volumes and snapshots, create peer connections and replication sets as described in [“Creating a peer connection” \(page 123\)](#), [“Creating a replication set from the Replications topic” \(page 126\)](#), and [“Creating a replication set from the Volumes topic” \(page 105\)](#).

Using the interface

Web browser requirements and setup

- Use Mozilla Firefox 57 and newer, Google Chrome 57 and newer, Microsoft Internet Explorer 10 and 11, or Apple Safari 10.1 and newer.

NOTE: If you use the Microsoft Edge browser that ships with Windows 10, you will be unable to view help content.

- To see the help window, you must 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):
 - For Internet Explorer, set the browser’s local-intranet security option to medium or medium-low.
 - Verify that the browser is set to allow cookies at least for the IP addresses of the storage-system network ports.
 - For Internet Explorer, add each controller’s network IP address as a trusted site.
 - If the SMU is configured to use HTTPS, ensure that Internet Explorer is set to use TLS 1.2.

Areas of the interface



The main areas of the SMU interface are the banner, topic tabs, topic pane, and footer, as represented by the following table. For information about a topic tab or an item in the banner or footer, click its link in the table.

The topic pane shows information that relates to the selected topic tab. This area also contains an Action menu that provides access to configuration, provisioning, and other actions. The contents of the Action menu are determined by the user’s role, the selected topic, and what (if anything) is selected in the topic pane.

Table 1 Areas of the SMU interface

Banner:	Product ID	System panel (page 139)	Connection panel (page 140)	Date/time panel (page 140)	User panel (page 141)	Sign Out button (page 15)	Help button (page 13)
Topic tabs:	Home (page 37)	Topic pane					
	System (page 61)						
	Hosts (page 80)						
	Pools (page 86)						
	Volumes (page 96)						
	Mapping (page 111)						
	Replication (page 116)						
	Performance (page 135)						
Footer:	Health panel (page 141)	Event panel (page 142)	Capacity panel (page 143)	Host I/O panel (page 144)	Tier I/O panel (page 144)	I/O Workload panel (page 144)	Activity panel (page 144)

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.
- A red asterisk (*) identifies a required setting.
- As you set options in action panels, the SMU informs you whether a value is invalid or a required option is not set. If the **Apply** or **OK** button remains inactive after you set all required options, either press **Tab** or click in an empty area of the panel to activate the button.
- If an action panel has an **Apply** button and an **OK** button, click **Apply** to apply any changes and keep the panel open or click **OK** to apply any changes and close the panel. After clicking **Apply**, you can click **Close** to close the panel without losing changes already applied.
- You can move an action panel or a confirmation panel by dragging its top border.
- 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.
- 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.
- If you start to perform an action in a panel (such as adding a new entry to a table) and then select an item or button that interrupts the action, a confirmation panel will ask if you want to navigate away and lose any changes made. If you want to continue performing the original action, click **No**. If you want to stop performing the original action, click **Yes**.
- In the banner or footer,  or  indicates that a panel has a menu. Click anywhere in the panel to display the menu.
- Right-clicking on a row in a topic table displays that topic's Action menu. This provides an additional and faster method for more experienced users to access the menu items. Hovering over a disabled menu item provides a tool tip indicating why the item is disabled.



Tips for using tables

Items such as initiators, hosts, volumes, and mappings are listed in tables. Use the following methods singly or together to quickly locate items that you want to work with.

Selecting items

- To select an item, click in its row.
- To select a range of adjacent items, click the first item in the range and **Shift+click** the last item in the range.
- To select or deselect one or more items, **Ctrl+click** each one.

Sorting items

To sort items by a specific column, click the column heading to reorder items from low to high (). Click again to reorder items from high to low ().

To sort items by multiple columns

1. In the first column to sort by, click its heading once or twice to reorder items.
2. In the second column to sort by, **Shift+click** its heading once or twice to reorder items. If you **Shift+click** a third time, the column is deselected.
3. Continue for each additional column to sort by.


Using filters to find items with specified text

To filter a multicolumn table, in the filter field above the table, enter the text to find. As you type, only items that contain the specified text remain shown. Filters are not case sensitive.

To use a column filter

1. In the column heading click the filter icon (). The filter menu appears.

2. Do one of the following:

- In the filter field, enter the text to find. As you type, only items that contain the specified text remain shown. Because a filter is active, the icon changes (). Previous search terms are listed below the field. Previous search terms that match displayed values are shown in bold.
- If the filter list has an entry for the text you want to find, select that entry.
- To show all items in the column, click the filter icon and select **All**.

To clear all filters and show all items, click **Clear Filters**.

Limiting the number of items shown

To show a specific number of items at a time in a multicolumn table, select a value from the **Show** menu. If more items exist, you can page through them by using the following buttons:

- > Show next set of items.
- > Reached end of list.
- < Show previous set of items.
- < Reached start of list.

Exporting data to a CSV file

You can export initiator, host, volume, mapping, and replication data that is displayed in tables to a downloadable Comma Separated Values (CSV) file that can be viewed in a spreadsheet for further analysis. Data can be exported for the entire table or for one or more selected rows, and it can be displayed in row format or column format. The exported CSV file contains all of the data in the table including information that is displayed in the hover panels.





To export table data to a CSV file

1. Select one or more rows of data to export from a table that has an Export to CSV button.
2. Click **Export to CSV**. The Export Data to CSV panel opens.
3. Click **All** to export all of the data within the selected table, or click **Selected** to export only selected files.
4. Click **Rows** to export the data in row format, or **Columns** to export the data in column format.
5. Click **OK**. The data is exported to a CSV file.

Tips for using help

- To display help for the content in the topic pane, click the help icon  in the banner.

NOTE: If you use the Microsoft Edge browser that ships with Windows 10, you will be unable to view help content.














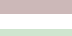

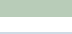
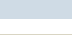
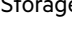

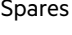

- In the help window, click the table of contents icon  to show or hide the Contents pane.
- As the context in the main panel is changed, the corresponding help topic is displayed in the help window. To prevent this automatic context-switching, click the pin icon . When a help window is pinned (), you can still browse to other topics within the window and you can open a new window. You cannot unpin a help window. You can only close it.
- If you have viewed more than one help topic, you can click the arrow icons to display the previous or next topic.
- To close the help window, click the close icon .

Color codes

The interface uses the following color codes to distinguish performance statistics and types of capacity utilization.





Home topic

Table 2 Home topic storage space color codes

Color	Meaning
System performance statistics	
	IOPS
	Data throughput (MB/s)
Capacity graph, bottom bar	
	System physical space available
	System physical space used by global spares
	System physical space used by virtual disk groups
Capacity graph, top bar	
	Virtual pool reserved space (RAID parity and metadata)
	Virtual pool allocated space
	Virtual pool unallocated space
Storage A/B, virtual capacity graph, bottom bar	
	Virtual pool usable space (excludes reserved space)
Storage A/B, virtual capacity graph, top bar	
	Virtual pool allocated space
	Virtual pool unallocated space
Storage A/B, virtual disk group utilization graph	
	Performance tier unallocated space
	Performance tier allocated space
	Standard tier unallocated space
	Standard tier allocated space
	Archive tier unallocated space
	Archive tier allocated space
Storage A/B, read cache utilization graph	
	Read cache unallocated space
	Read cache allocated space
Spares	
	Standard tier global spares
	Archive tier global spares

Create Virtual Volumes panel

Table 3 Create Virtual Volumes panel storage space color codes

Color	Meaning
Virtual capacity graph, top bar	
	Virtual pool allocated space
	Virtual pool unallocated space
	Virtual pool space that would be used by the volumes being created
Virtual capacity graph, bottom bar	
	Virtual pool usable space (excludes reserved space)

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. Disk drives usually show size in base 10. Memory (RAM and ROM) size is always shown in base 2. In the SMU, the base for entry and display of storage-space sizes can be set per user. When entering storage-space sizes only, either base-2 or base-10 units can be specified. Base 10 is the default. For more information about setting base types, see [“Managing users” \(page 42\)](#).

Table 4 Storage size representations in base 2 and base 10

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

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

Table 5 Decimal (radix) point character by locale

Language	Character	Examples
English, Chinese, Japanese, Korean	Period (.)	146.81 GB 3.0 Gbit/s
Dutch, French, German, Italian, Spanish	Comma (,)	146,81 GB 3,0 Gbit/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. Depending on how your browser treats this session identifier, you might be able to run multiple independent sessions simultaneously. For example, each instance of Internet Explorer can run a separate SMU session, but all instances of Firefox, Chrome, and Safari share the same SMU session.

To sign in

1. In the web browser address field, type `https://<IP address of a controller network port>` and press **Enter**. (Do not include a leading zero in an IP address. For example, enter 10.1.4.33 and not 10.1.4.033.) The SMU Sign In page is displayed. 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” \(page 50\)](#) or the `set protocols` CLI command.

2. On the sign-in page, enter the name and password of a configured user. The default user name and password for someone who can both monitor and manage the system is `manage` and `!manage`. The default user name and password for someone who can only monitor the system is `monitor` and `!monitor`.
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. Click **Sign In**. If the system is available, the Home page or the Welcome panel displays. Otherwise, a message indicates that the system is unavailable.

When you are ready to end your session, sign out as described below. Do not simply close the browser window.

To sign out

1. Click **Sign Out** near the top of the SMU window.
2. In the confirmation panel, click **Sign Out**.

System concepts

About 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 RAID sets that are grouped into a pool. Thus, a virtual volume as seen by a host represents a portion of storage in a pool. Multiple virtual volumes can be created in a pool, sharing its resources. This allows for a high level of flexibility, and the most efficient use of available physical resources.

Some advantages of using virtual storage are:

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

Virtual storage provides the foundation for data-management features such as thin provisioning on [page 24](#), automated tiered storage on [page 25](#), read cache on [page 21](#), and the quick rebuild feature on [page 29](#).

About the handling of linear storage

The controller modules in this product support virtual storage but not linear storage. If you upgraded a system that supported linear storage, or inserted disks that contain linear disk groups from a different system, read this topic to understand how that data will be handled.

When the system finds linear disk groups, they are quarantined. You cannot perform read or write operations to any volumes in those disk groups; you can only view and delete them. No configuration actions—including but not limited to mapping, replication, and creating volumes—will be allowed on a linear disk group.

The system manages linear objects as follows:

- Linear disk groups display in the Pools topic and are assigned the status of QTUN (quarantined, unsupported). Unsupported disk groups may not be dequarantined and any attempt to do so will result in an error message. Deleting the disk group removes the following associated objects from the system: volumes, snapshots, snap pools, and linear replication schedules.
- Linear volumes and all snapshots associated with them display in the Volumes topic until their associated disk groups are deleted. Snap pools are not visible.
- Linear replication sets are not visible in the system but still remain associated with a disk group. You will be unable to create virtual replication sets until these objects, through their associated disk groups, are deleted from the system.
- Dedicated spares will remain designated as such until their associated linear disk group is deleted. At that time they will return to an available state.
- Global spares configured in a linear system will remain available.
- Linear mapped volumes may sometimes be unmapped while quarantined. Mappings are automatically removed when their associated disk group is deleted.

NOTE: If you still need access to the linear-provisioned data, place the disks in a system that supports linear disk groups. If you no longer need the linear-provisioned data, you can delete the quarantined disk groups.

NOTE: Tasks, schedules, and remote system definitions that were stored on the previous controllers will not be available in this system.

To delete quarantined linear disk groups

1. In the Pools topic, select the pool for the disk group(s) that you are deleting in the pools table. Then, select the linear disk group(s) in the Related Disk Groups table.
2. Select **Action > Remove Disk Groups**. The Remove Disk Groups panel opens.
3. Click **OK**.
4. Click **Yes** to continue. Otherwise, click **No**. If you clicked **Yes**, the disk group(s) and all associated objects including volumes, snap pools, snapshots, and replication artifacts are deleted, and the Related Disk Groups table is updated.

About disk groups

A *disk group* is an aggregation of disks of the same type, using a specific RAID level that is incorporated as a component of a pool, for the purpose of storing volume data. You can add virtual and read-cache disk groups to a pool.

All disks in a disk group must be the same type (SSD, enterprise SAS, or midline SAS). A disk group can contain different models of disks, and disks with different capacities and sector formats. 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 level. For example, the capacity of a disk group composed of one 500 GB disk and one 750 GB disk is equivalent to a disk group composed of two 500 GB disks. To maximize capacity, use disks of similar size.

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).

You can provision storage by adding a disk group to a pool. Volumes then can be created in the pool.

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 groups

A virtual disk group requires the specification of a set of disks, RAID level, disk group type, pool target (A or B), and a name. 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. Virtual disk groups that contain SSDs can only be created with a Performance Tier license when other HDD disk groups exist in the system. If the system contains only SSDs, then virtual disk groups can be created without a Performance Tier license. This restriction does not apply to read-cache disk groups.

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.
-

When a virtual disk group is removed that contains active volume data, that volume data will drain (or be moved) to other disk group members within the pool (if they exist). Disk groups should only be removed when all volume data can cleanly be drained from the disk group. When the last disk group is removed, the pool ceases to exist and will be deleted from the system automatically.

NOTE: If the last disk group contains data, a warning will display prompting you to confirm removing the disk group.

The RAID level for a virtual disk group must be fault tolerant. The supported RAID levels for virtual disk groups are: RAID 1, RAID 5, RAID 6, RAID 10. If RAID 10 is specified, the disk group must have at least two sub-groups.

Read-cache disk groups


A read-cache disk group is a special type of a virtual disk group that is used to cache virtual pages to improve read performance. Read cache does not add to the overall capacity of the pool to which it has been added. You can add or remove it from the pool without any adverse effect on the volumes and their data for the pool, other than to impact the read-access performance.

If your system uses SSDs, you can create read-cache disk groups for virtual pools if you do not have any virtual disk groups for the pool that are comprised of SSDs (virtual pools cannot contain both read-cache and a Performance tier).

Only a single read-cache disk group may exist within a pool. Increasing the size of read cache within a pool requires the user to remove the read-cache disk group, and then re-add a larger read-cache disk group. It is possible to have a read-cache disk group that consists of one or two disks with a non-fault tolerant RAID level. For more information on read cache, see [“About SSD read cache” \(page 21\)](#).

About RAID levels

The RAID 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 RAID 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.

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

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)

NOTE: You can only create RAID-1, RAID-5, RAID-6, and RAID-10 virtual disk groups.

Table 6 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

Table 7 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
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

Table 8 Number of disks per RAID level to optimize virtual disk group performance

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

About SSDs

The use of SSDs (solid-state drives) can greatly enhance the performance of a system. Since the SSDs do not have moving parts, data that is random in nature can be accessed much faster. If you have the Performance Tier license, you can use SSDs for virtual disk groups. When combined with virtual disk groups that consist of other classes of disks, improved read and write performance is possible through automated tiered storage. Alternatively, you can use one or two SSDs in read-cache disk groups to increase read performance for pools without a Performance tier. The application workload of a system determines the percentage of SSDs of the total disk capacity that is needed for best performance.

For more information about automated tiered storage, see [“About automated tiered storage” \(page 25\)](#). For more information about read-cache disk groups, see [“Read-cache disk groups” \(page 18\)](#). For information about using SSDs in all disk groups, see [“All-flash array” \(page 20\)](#).

The rules for using SSDs and spinning disks are:

- If the first disk group is provisioned with SSDs and the system does not have the Performance Tier license installed, then the system will expect to be provisioned as an all-flash array and allow only SSDs to be used in all other disk groups. For more information on All Flash Array, see [“All-flash array” \(page 20\)](#).
- If the first disk group is provisioned with spinning disks and does not have a Performance Tier license installed, then the system can only be provisioned to use spinning disks in virtual disk groups and use SSDs as read cache.
- If the first disk group is provisioned with spinning disks and has a Performance Tier license installed, then the system can be provisioned to use spinning disks in virtual disk groups and use SSDs either in virtual disk groups or as read cache.

The application workload of a system determines the percentage of SSDs of the total disk capacity that is needed for best performance.

Gauging the percentage of life remaining for SSDs

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

You can view the value of the SSD Life Left property through the Disk Information panel. In the front view of the enclosure in the System topic, hover the cursor over any disk to view its properties. You can also view the Disk Information panel through the Pools topic. Select the pool for the disk group in the pools table, select the disk group in the Related Disk Groups table, and then hover the cursor over the disk in the Related Disks table.

All-flash array

An all-flash array is a storage system in which all disk groups use only SSDs, providing the ability to have a homogeneous SSD-only configuration. Tiering is not supported for an all-flash array. If a system includes disk groups with spinning disks, the disk groups must be removed before all-flash array can be used. For information about the rules for using SSDs and spinning disks, see [“About SSDs” \(page 20\)](#).

Internal disk management

SSDs use multiple algorithms to manage SSD endurance features. These include wear leveling, and over-provisioning to minimize write amplification.

Wear leveling

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

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

Overprovisioning

The write amplification factor of an SSD is defined as the ratio of the amount of data actually written by the SSD to the amount of host/user data requested to be written. This is used to account for the user data and activities like wear leveling. This affects wear leveling calculations and is influenced by the characteristics of data written to and read from SSDs. Data that is written in sequential LBAs that are aligned on 4KB boundaries results in the best write amplification factor. The worst write amplification factor typically occurs for randomly written LBAs of transfer sizes that are less than 4KB and that originate on LBA's that are not on 4KB boundaries. Try to align your data on 4KB boundaries.

Data retention

Data retention is another major characteristic of SSDs that all SSD algorithms take into account while running. While powered up, the data retention of SSD cells are monitored and rewritten if the cell levels decay to an unexpected level. Data retention when the drive is powered off is affected by Program and Erase (PE) cycles and the temperature of the drive when stored.

Drive Writes per Day (DWD)

DWD or DWPD refers to Drive Writes Per Day. Disk vendors rate SSD endurance by how many writes can occur over the lifetime of an SSD. As lower-cost SSDs that support fewer drive writes per day become available, the cost/benefit analysis of which SSDs to use is highly dependent on your applications and I/O workload, as is the ratio of SSDs to conventional drives. In some environments, a ratio of 10% SSDs to 90% conventional drives, when combined with HPE MSA real-time tiering, can yield dramatic performance improvements.

Since HPE MSA real-time tiering automatically moves “hot” data to SSDs and less-used “cool” data to conventional disks, applications and environments that require mission-critical movement of frequently accessed “hot” data might dictate a higher ratio of SSDs to conventional disks.

About SSD read cache

Unlike tiering, where a single copy of specific blocks of data resides in either spinning disks or SSDs, the Read Flash Cache (RFC) feature uses one SSD read-cache disk group per pool as a read cache for “hot” pages only. Each read-cache disk group consists of one or two SSDs with a maximum usable capacity of 4TB. A separate copy of the data is also kept in spinning disks. Read-cache contents are lost when a controller restart or failover occurs. Taken together, these attributes have several advantages:

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

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

For more information on SSDs, see [“About SSDs” \(page 20\)](#).

About spares

Spare disks are unused disks in your system that you designate to automatically replace a failed disk, restoring fault tolerance to disk groups in the system. 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.

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

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.

About 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 that is incorporated as a component of a pool, that stores volume data. For virtual pools, when volumes are added to a pool the data is distributed across the pool's disk groups.

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

You can provision disks into disk groups. For information about how provisioning disks works, see [“Adding a disk group” \(page 88\)](#).

Virtual pools and disk groups

The volumes within a virtual pool are allocated virtually (separated into fixed size pages, with each page allocated randomly from somewhere 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).

If you would like to create a virtual pool that is larger than 512 TiB on each controller, you can enable the large pools feature by using the `large-pools` parameter of the `set advanced-settings` CLI command. When the large pools feature is disabled, which is the default, the maximum size for a virtual pool is 512 TiB and the maximum number of volumes per snapshot tree is 255 (base volume plus 254 snapshots). Enabling the large pools feature will increase the maximum size for a virtual pool to 1024 TiB (1PiB) and decrease the maximum number of volumes per snapshot tree to 9 (base volume plus 8 snapshots). The maximum number of volumes per snapshot will decrease to fewer than 9 if more than 3 replication sets are defined for volumes in the snapshot tree. For more information about the `large-pools` parameter of the `set advanced-settings` CLI command, see the CLI documentation.

NOTE: The physical capacity limit for a virtual pool is 512 TiB. When overcommit is enabled, the logical capacity limit is 1 PiB.

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

If a system has at least one SSD, each virtual pool can also have a read-cache disk group. Unlike the other disk group types, read-cache disk groups are used internally by the system to improve read performance and do not increase the available capacity of the pool.

About volumes and volume groups

A volume is a logical subdivision of a virtual pool and can be mapped to host-based applications. A mapped 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 mapping, see [“About volume mapping” \(page 26\)](#).

Virtual volumes

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

Because virtual volumes and snapshots share the same underlying structure, it is possible to create snapshots of other snapshots, not just of volumes, creating a snapshot tree.

A maximum of 1024 virtual volumes can exist per system.

Volume groups

You can group a maximum of 1024 volumes (standard volumes, snapshots, or both) into a volume group. Doing so enables you to perform mapping operations for all volumes in a group at once, instead of for each volume individually. A volume can be a member of only one group. All volumes in a group must be in the same virtual pool. A volume group cannot have the same name as another volume group, but can have the same name as any volume. A maximum of 256 volume groups can exist per system. If a volume group is being replicated, the maximum number of volumes that can exist in the group is 16.

About 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.

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.

When modifying a volume you can change its write-back cache setting. *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 mirrors all of the data from one controller module cache to the other. 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 [“Changing system cache settings” \(page 75\)](#).

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 corruption or loss. Before changing this setting, quiesce I/O from all initiators.

You can also change the optimization mode.

- **Standard.** This controller cache mode of operation is optimized for sequential and random I/O and is the optimization of choice for most workloads. In this mode, the cache is kept coherent with the partner controller. This mode gives you high performance and high redundancy. This is the default.
- **No-mirror.** In this mode of operation, the controller cache performs the same as the standard mode with the exception that the cache metadata is not mirrored to the partner. While this improves the response time of write I/O, it comes at the cost of redundancy. If this option is used, the user can expect higher write performance but is exposed to data loss if a controller fails.

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 NRAID and RAID-1 disk groups internally as if they have a stripe size of 512 KB, even though they are not striped.
- Specific size options let you select an amount of data for all accesses.
- 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.

About thin provisioning

Thin provisioning is a virtual storage feature that allows a system administrator to overcommit physical storage resources. This allows the host system to operate as though it has more storage available than is actually allocated to it. When physical resources fill up, the administrator can add physical storage by adding additional disk groups, on demand.

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

For example, contrast the methods for creating a volume for Microsoft Exchange Server data:

- Typically, administrators create a storage-side volume for Exchange and map that volume with an assigned LUN to hosts, and then create a Microsoft Windows volume for that LUN. Each volume has a fixed size. There are ways to increase the size of a storage-side volume and its associated Windows volume, but they are often cumbersome. The administrator must make a trade-off between initial disk costs and a volume size that provides capacity for future growth.
- With thin provisioning, the administrator can create a very large volume, up to the maximum size allowed by Windows. The administrator can begin with only a small number of disks, and add more as physical storage needs grow. The process of expanding the Windows volume is eliminated.

NOTE: For a thin-provisioned volume mapped to a host, when data is deleted from the volume not all of the pages (space) associated with that data will be deallocated (released). This is especially true for smaller files. To deallocate the pages, in Windows, select the mapped volume and do either of the following:

- Perform a quick format.
 - View its properties, select the **Tools** tab, and under **Defragmentation**, click **Optimize**.
-

About automated tiered storage

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

- Frequently accessed, “hot” data can move to disks with higher performance.
- Infrequently accessed, “cool” data can move to disks with lower performance and lower costs.

Each virtual disk group, depending on the type of disks it uses, is automatically assigned to one of the following tiers:

- **Performance**—This highest tier uses SSDs, which provide the best performance but also the highest cost. For more information on SSDs, see [“About SSDs” \(page 20\)](#).
- **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 critical (CRIT), the system will automatically drain data from that disk group to disk groups using spinning disks in other tiers providing that they can contain the data on the degraded disk group. This occurs because similar wear across the SSDs is likely, so more failures may be imminent.

If a system only has one class of disk, no tiering occurs. However, automated tiered storage rebalancing happens when adding or removing a disk group in a different tier.

Volume tier affinity feature

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 available performing tiers first and only uses the Archive tier when space is exhausted in the other tiers. Volume data will swap into higher performing tiers based on frequency of access and tier space availability. This is the default.
- **Archive**—This setting prioritizes the volume data to the least 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 volume data to the higher performing tiers. If no space is available, lower performing tier space is used. Performance affinity volume data will swap into higher tiers based upon frequency of access or when space is made available.

About 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 FC HBA) in a server.

The controllers automatically discover initiators that have sent an `inquiry` command or a `report luns` command to the storage system, which typically happens when a host boots up or rescans for devices. When the command is received, the system saves the initiator ID. You can also manually create entries for initiators. For example, you might want to define an initiator before a controller port is physically connected through a switch to a host.

You can assign a nickname to an initiator to make it easy to recognize for volume mapping. For a named initiator, you can also select a profile specific to the operating system for that initiator. A maximum of 512 names can be assigned.

For ease of management, you can group 1–128 initiators that represent a server into a host. Further, you can group 1–256 hosts into a host group. Doing so enables you to perform mapping operations for all initiators in a host, or all initiators and hosts in a group, instead of for each initiator or host individually. An initiator must have a nickname to be added to a host, and 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.

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. Steps involved in enabling CHAP include:

- Decide on host node names (identifiers) and secrets. The host node name is its IQN. A secret must have 12–16 characters.
- Define CHAP entries in the storage system.
- Enable CHAP on the storage system. 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 CHAP secret in the host iSCSI initiator.
- Establish a new connection to the storage system using CHAP. The host should be displayable by the system, as well as the ports through which connections were made.

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. This information 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 move to the full feature phase.

About volume mapping

Mappings between a volume and one or more initiators, hosts, or host groups (hereafter called “hosts”) enable the hosts to view and access the volume. There are two types of maps that can be created: default maps and explicit maps. Default maps enable all hosts to see the volume using a specified LUN and access permissions. Default mapping applies to any host that has not been explicitly mapped using different settings. Explicit maps override a volume's default map for specific hosts.

The advantage of a default mapping is that all connected hosts can discover the volume with no additional work by the administrator. The disadvantage is that all connected hosts can discover the volume with no restrictions. Therefore, this process is not recommended for specialized volumes that require restricted access. Also, to avoid multiple hosts mounting the volume and causing corruption, the hosts 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 corruption. To control access by specific hosts, you can create an explicit mapping. An explicit mapping can use a different access mode, LUN, and port settings to allow or prevent access by a host to a volume. If there is a default mapping, the explicit mapping overrides it.

When a volume is created, it is not mapped by default. You can create default or explicit mappings for it.

You can change the default mapping of a volume, and create, modify, or delete explicit mappings. A mapping can specify read-write, read-only, or no access through one or more controller host ports to a volume. When a mapping specifies no access, the volume is masked.

For example, a payroll volume could be mapped with read-write access for the Human Resources host and be masked for all other hosts. An engineering volume could be mapped with read-write access for the Engineering host and read-only access for other departments' hosts.

A LUN identifies a mapped volume to a host. Both controllers share a set of LUNs, and any unused LUN can be assigned to a mapping. However, each LUN is generally only used once as a default LUN. For example, if LUN 5 is the default for Volume1, no other volume in the storage system can use LUN 5 on the same port as its default LUN. For explicit mappings, the rules differ: LUNs used in default mappings can be reused in explicit mappings for other volumes and other hosts.

 **TIP:** When an explicit mapping is deleted, the volume's default mapping takes effect. Though default mappings can be used for specific installations, using explicit mappings with hosts and host groups is recommended for most installations.

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 in 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 disk group ownership. When ULP is in use, 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.

About 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 cannot be expanded.

A base of 64 snapshots is included with all systems without an additional license. With a license, you can create up to the maximum number of snapshots for your product. When you reach the maximum base number of snapshots, before you can create a new snapshot you must either delete an existing snapshot or purchase and install a license that increases the maximum number of snapshots. The system treats a snapshot like any other volume. The snapshot can be mapped to hosts with read-only access, read-write access, or no access, depending on the purpose of the snapshot.

Snapshots use the rollback feature which replaces the data of a source volume or snapshot with the data of a snapshot that was created from it. Snapshots also use the reset snapshot feature that enables you to replace the data in a snapshot with the current data in the source volume or snapshot. When you reset a snapshot, the snapshot name and mappings are not changed.

The `set snapshot-space` CLI command enables you to set the percent of the pool that can be used for snapshots (the snapshot space). Optionally, you can specify a limit policy to enact when the snapshot space reaches the percentage. You can set the policy to either notify you via the event log 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.

Virtual snapshots

Creating snapshots is a fast and efficient process that merely consists of pointing to the same data to which the source volume or snapshot points. (Since snapshots reference volumes, they take up no space unless they or the source volume or source snapshot is modified.) Space does not have to be reserved for snapshots because all space in the pool is available for them. 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.

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

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

delete existing snapshots from the tree. You can only delete snapshots that do not have any child snapshots. You cannot expand the base volume of a snapshot tree or 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 virtual snapshots are copies of a point in time, a modified snapshot cannot be reverted. If you want a virtual snapshot to provide the capability to “revert” the contents of the source volume or snapshot to when the snapshot was created, create a snapshot for this purpose and archive it so you do not change the contents.

For 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.

About copying volumes

The volume copy feature enables you to copy a virtual base volume and snapshot to a new volume. It creates a complete “physical” copy of a base volume or a snapshot within a storage system. It is an exact copy 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 an 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.

For more information about creating a copy of a virtual base volume or snapshot, see [“Copying a volume or snapshot” \(page 100\)](#).

About reconstruction

If one or more disks fail in a disk group and spares of the appropriate size (same or larger) and type (same as the failed disks) are available, the storage system automatically uses the spares 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 no spares are available, reconstruction does not start automatically. To start reconstruction manually, replace each failed disk and designate each replacement disk as a spare. If you have configured the dynamic spares feature through the CLI, reconstruction will automatically start for disk groups. With dynamic spares enabled, if a disk fails and you replace it with a compatible disk, the storage system rescans the bus, finds the new disk, automatically designates it a spare, and starts reconstructing the disk group (as described in [“About spares” \(page 22\)](#)).

For virtual storage, reconstruction of all disk groups uses a quick-rebuild feature. For more information on quick rebuild, see [“About quick rebuild” \(page 29\)](#).

When a disk fails, its fault LED illuminates amber. When a spare is used as a reconstruction target, its activity LED blinks green. During reconstruction, the fault LED and activity LEDs for all disks in the disk group blink. For descriptions of LED states, see the User Guide.

NOTE: Reconstruction can take hours or days to complete, depending on the disk group RAID level and size, disk speed, utility priority, 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.

About 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 only 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 RAID rebuild. Data stripes that have not been allocated to user data are scrubbed in the background, using a lightweight process that allows future data allocations to be more efficient.

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

About performance statistics

You can view current or historical performance statistics for components of the storage system.

Current performance statistics for disks, disk groups, pools, tiers, host ports, controllers, and volumes are displayed in tabular format. Current statistics show the current performance and are sampled immediately upon request.

Historical performance statistics for disks, pools, and tiers are displayed in graphs for ease of analysis. Historical statistics focus on disk workload. You can view historical statistics to determine whether I/O is balanced across pools and to identify disks that are experiencing errors or are performing poorly.

The system samples historical statistics for disks every quarter hour and retains these samples for 6 months. It samples statistics for pools and tiers every 5 minutes and retains this data for one week but does not persist it across failover or power cycling. By default, the graphs show the latest 100 data samples, but you can specify either a time range of samples to display or a count of samples to display. The graphs can show a maximum of 100 samples.

If you specify a time range of samples to display, the system determines whether the number of samples in the time range exceeds the number of samples that can be displayed (100), requiring aggregation. To determine this, the system divides the number of samples in the specified time range by 100, giving a quotient and a remainder. If the quotient is 1, the 100 newest samples will be displayed. If the quotient exceeds 1, each “quotient” number of newest samples will be aggregated into one sample for display. The remainder is the number of oldest samples that will be excluded from display.

- Example 1: A 1-hour range includes 4 samples. 4 is less than 100 so all 4 samples are displayed.
- Example 2: A 30-hour range includes 120 samples. 120 divided by 100 gives a quotient of 1 and a remainder of 20. Therefore, the newest 100 samples will be displayed and the oldest 20 samples will be excluded.
- Example 3: A 60-hour range includes 240 samples. 240 divided by 100 gives a quotient of 2 and a remainder of 40. Therefore, each two newest samples will be aggregated into one sample for display and the oldest 40 samples will be excluded.

If aggregation is required, the system calculates values for the aggregated samples. For a count statistic (total data transferred, data read, data written, total I/Os, number of reads, number of writes), the samples' values are added to produce the value of the aggregated sample. For a rate statistic (total data throughput, read throughput, write throughput, total IOPS, read IOPS, write IOPS), the samples' values are added and then are divided by their combined interval. The base unit for data throughput is bytes per second.

- Example 1: Two samples' number-of-reads values must be aggregated into one sample. If the value for sample 1 is 1060 and the value for sample 2 is 2000 then the value of the aggregated sample is 3060.
- Example 2: Continuing from example 1, each sample's interval is 900 seconds so their combined interval is 1800 seconds. Their aggregate read-IOPS value is their aggregate number of reads (3060) divided by their combined interval (1800 seconds), which is 1.7.

You can export historical performance statistics in CSV format to a file on the network for import into a spreadsheet or other application. You can also reset current or historical statistics, which clears the retained data and continues to gather new samples.

For more information about performance statistics, see [“Viewing performance statistics” \(page 135\)](#), [“Updating historical statistics” \(page 137\)](#), [“Exporting historical performance statistics” \(page 137\)](#), and [“Resetting performance statistics” \(page 138\)](#).

About firmware update

Controller modules, expansion modules, and disk drives contain firmware that operate them. As newer firmware versions become available, they may be installed at the factory or they may be installed by storage-system administrators at customer sites. For a dual-controller system, the following firmware-update scenarios are supported:

- The administrator installs a new firmware version in one controller and wants that version to be transferred to the partner controller.
- In a system that has been qualified with a specific firmware version, the administrator replaces one controller module and wants the firmware version in the remaining controller to be transferred to the new controller (which might contain older or newer firmware).

When a controller module is installed into an enclosure at the factory, the enclosure midplane serial number and firmware-update timestamp are recorded for each firmware component in controller flash memory, and will not be erased when the configuration is changed or is reset to defaults. These two pieces of data are not present in controller modules that are not factory-installed and are used as replacements.

Updating controller firmware with the Partner Firmware Update (PFU) option enabled will ensure that the same firmware version is installed in both controller modules. 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.
- If the firmware in only one controller has the proper midplane serial number then the firmware, midplane serial number, and attributes of that controller are transferred to the partner controller. Subsequently, the firmware update behavior for both controllers depends on the system settings.
- If the firmware in both controllers has the proper midplane serial number then the firmware having the latest firmware-update timestamp is transferred to the partner controller.
- If the firmware in neither controller has the proper midplane serial number, then the firmware version in controller A is transferred to controller B.

For information about the procedures to update firmware in controller modules, expansion modules, and disk drives, see [“Updating firmware” \(page 66\)](#). That topic also describes how to use the activity progress interface to view detailed information about the progress of a firmware-update operation.

About 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. The managed logs feature allows log data to be transferred to a log-collection system, and store it for later retrieval before any data is lost. The *log-collection system* is a host computer that is designated to receive the log data transferred from the storage system. The transfer does not remove any data from the logs in the storage system. This feature is disabled by default.

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

- In push mode, when log data has accumulated to a significant size, the storage system sends notifications with attached log files via email to the log-collection system. The notification will specify the storage-system name, location, contact, and IP address, and will contain a single log segment in a compressed zip file. The log segment will be uniquely named to indicate the log-file type, 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 notifications via email, SMI-S, or SNMP to the log-collection system, which can then use FTP or SFTP to transfer the appropriate logs from the storage system. The notification will specify the storage-system name, location, contact, and IP address and the log-file type (region) that needs to be transferred.

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

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

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

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

Following the transfer of 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 its partner will send the logs from both controllers.

Alternative methods for obtaining log data are to use the Save Logs action in the SMU or the `get_logs` command in the FTP or SFTP interface. These methods will transfer the entire contents of a log file without changing its capacity-status level. Use of Save Logs or `get_logs` is expected as part of providing information for a technical support request. For information about using the Save Logs action, see [“Saving log data to a file” \(page 141\)](#). For information about using the FTP or SFTP interface, see [“Using FTP and SFTP” \(page 163\)](#).

About LDAP

You can configure the storage system to use external Lightweight Directory Access Protocol (LDAP) services provided from Windows 2016 or 2012 R2 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 tab of the Manage Users panel 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, see [“Managing users” \(page 42\)](#). 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 or 2012 R2 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 tab in the Manage Users panel 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.

LDAP Integration

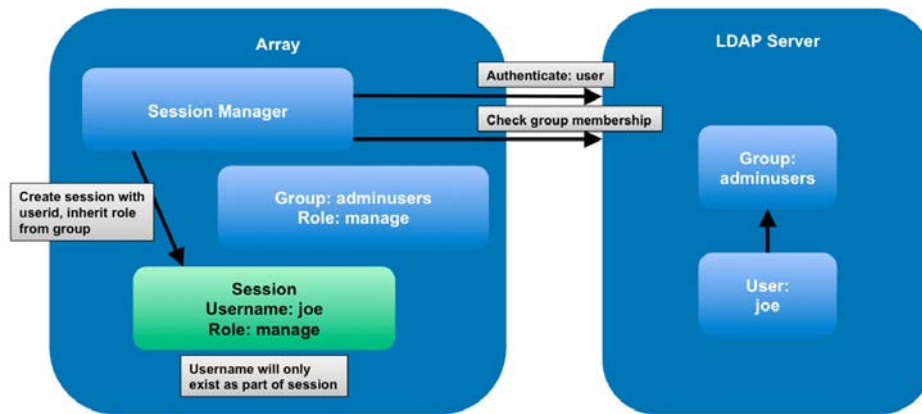


Figure 1 Replication process for initial replication

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 include the storage size base, precision, and units; locale; temperature-scale; and timeout. 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).

Audit logs

User login, logout, and actions through all interfaces for both local and LDAP users will be recorded in an audit log. 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, FTP/SFTP, and SMI-S protocols, but will not contain specific value changes, such as old and new settings.

Audit logs record host IP information for all interfaces. Audit logs also record SNMP SET requests.

Each controller maintains its own audit log. Each audit log can contain up to 2MB of data, after which it will wrap.

Audit log data will persist after the CLI `restore defaults` command is run. Audit log data 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.

Protocols and services

Before enabling the LDAP feature, unsecure protocols and services (Telnet, HTTP, FTP, unsecure SMI-S, debug, and activity progress reporting via HTTP) must be disabled. When the LDAP feature is enabled, only secure protocols and services (SSH, HTTPS, SFTP, and secure SMI-S) can be enabled. The LDAP feature must be disabled before unsecure protocols can be re-enabled. HTTPS, SSH, SFTP, and SMI-S 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 Windows 2016 or 2012 R2. The server must allow basic authentication using an LDAP over SSL (LDAPS) interface; that is, a TLS v1.2 connection.

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 MSA 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 time-out 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 login to the storage system either through the SMU or the CLI, the group from the 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 User 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. 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 organizational unit.
2. Check the **Enable LDAP Configuration** check box.
3. In the Server field, enter `10.235.217.52`.
4. In the Port field, enter `636`.
5. In the Alt-Server field, enter `10.235.217.51`.
6. In the Alt-Port field, enter `636`.
7. In the User search-base field, enter `ou=colo,dc=bigco2,dc=com,dc=local`.
8. Create an LDAP user group named `ArrayAdmins` (matching the group name on the LDAP server) with the manage role and access to the SMU and CLI interfaces.
9. In the User Group Name field, enter `ArrayAdmins`.
10. For Roles, check the **Manage** and **Monitor** checkboxes.
11. For Interfaces, check the **WBI** and **CLI** checkboxes.
12. Select the **Apply** button. When user `alice` attempts an SSH login to the storage system, the system connects to the configured LDAP server using the supplied credentials to perform authentication.

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

- Email-address format. For example:
`ssh alice@bigoc2.com.local@10.235.212.161`
- Domain\username format. For example:
`ssh bigco2\\alice@10.235.212.161`

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

NOTE: By default when creating a new user object in Windows Server 2016 or 2012 R2, 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.

About replicating virtual volumes

Replication 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.

For information about replication for virtual storage, see [“Working in the Replications topic” \(page 116\)](#).

Replicating between MSA 1050/2050 and MSA 1040/2040 systems

You can create and modify peer connections and replicate virtual volumes between a system with MSA 1050/2050 controllers and a system with MSA 1040/2040 controllers. However, some functionality will be limited. Read this topic to understand how feature compatibility is managed.

Rules for creating and modifying a peer connection

- You can create or modify a peer connection from an MSA 1050/2050 system to an MSA 1040/2040 system, but not from an MSA 1040/2040 system to an MSA 1050/2050 system. The system creating or modifying the peer connection must be an MSA 1050/2050.
- When creating or modifying a peer connection from an MSA 1050/2050 system to an MSA 1040/2040 system, the password for the MSA 1040/2040 system (the remote system relative to where the command is being issued) will not be checked and authorization will always succeed.
- You cannot create an FC peer connection between an MSA 1050/2050 system and an MSA 1040/2040 system.
- You can create an iSCSI peer connection between an MSA 1050/2050 system and an MSA 1040/2040 system.

Rules for using replication queue policy

- For replication from an MSA 1050/2050 system to an MSA 1040/2040 system the queue policy is set to `Discard`. Any CLI requests to change the policy will be rejected.
- For replication between MSA 1050/2050 systems the queue policy will default to `Queue Latest`.
- If the replication set was created on an MSA 1040/2040 system, and the disks of both the primary and secondary system are moved into an MSA 1050/2050 system, then you can change the replication set queue policy from `Discard` to `Queue Latest`.
- Moving disks from a MSA 1050/2050 system to a MSA 1040/2040 system will change any existing queue policy to `Discard` regardless of its previous setting.

Rules for using replication snapshot history

- This feature is only available for replication between MSA 1050/2050 systems.
- Snapshots of volume group replications are not supported.
- Replication sets are retained when moving disks from MSA 1040/2040 systems to an MSA 1050/2050 system. After the upgrade, existing replication sets can use the Replication Snapshot History feature.
- Snapshot history data is retained when moving disks from MSA 1050/2050 systems to MSA 1040/2040 systems. However, the replication snapshot history feature will no longer be available.

About the Full Disk Encryption feature (for MSA 2050 only)

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

The ability to secure a disk and system relies on passphrases and lock keys. A passphrase is a user-created password that allows users to manage lock keys. A lock key is generated by the system and manages the encryption and decryption of data on the disks. A lock key is persisted on the storage system and is not available outside the storage system.

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.

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

Secured disks and systems can be repurposed. Repurposing a disk changes the encryption key on the disk, effectively erasing all data on the disk and unsecuring the system and disks. Repurpose a disk only if you no longer need the data on the disk.

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. For information on setting up FDE and modifying FDE options, see [“Changing FDE settings \(for MSA 2050 only\)” \(page 70\)](#).

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 disk channels” \(page 64\)](#).

About 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 section presents some considerations concerning data protection.

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

If the controller fails while in write-back mode, unwritten cache data likely exists. The same is true if the controller enclosure 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, please 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 [“About volume cache options” \(page 23\)](#). For more information about changing cache settings for a volume, see [“Modifying a volume” \(page 100\)](#). For more information about changing system cache settings, see [“Changing system cache settings” \(page 75\)](#).

About SAS cabling (for MSA 1050 only)

For systems with a 2-port SAS controller module, host ports can be configured through the SMU or CLI to use fan-out cables or standard cables. A standard cable can connect one port on a SAS host to one controller port, using four PHY lanes per port. A fan-out cable can connect one port on each of two SAS hosts to one controller port, using two PHY lanes per port. Using fan-out cables instead of standard cables doubles the number of hosts that can be attached to a single system. It will also halve the maximum bandwidth available to each host, but overall bandwidth available to all hosts is unchanged. Configuration must be the same for all ports on both controllers, so a mix of standard cables and fan-out cables cannot be used on one system. Use of fan-out cables is enabled by default.

Once you have switched the configuration through the firmware, you can disconnect the existing cables and switch to the other type of cables. For information on how to connect and disconnect cables, refer to your product's User Guide.

If you connect a cable that does not match the cable type for the configuration, an event will be logged that indicates a mismatch has occurred. Also, while I/O will occur, half of the PHY lanes for each port will be disabled. The Ports hover panel accessed through the Home topic will reflect that the port is in a degraded state. If a cable mismatch occurs, change the port mode of the system using the Host Ports Settings panel or connect cables of the appropriate type for the configuration.

For more information on checking port properties through the Home topic, see [“Port information” \(page 38\)](#).

When configuring the host-interface settings for a 2-port SAS controller module, the current link speed, number of PHY lanes expected for the SAS port, and number of PHY lanes active for each SAS port are displayed. The number of ports that appear depends on the configuration. Changing the host-interface settings interrupts I/O and restarts the storage controllers. For more information on how to configure host ports for use with SAS fan-out cables, see [“Resetting host ports” \(page 64\)](#).

2 Working in the Home topic

The Home topic provides options to set up and configure your system and manage tasks, and displays an overview of the storage managed by the system. The content that displays depends on the completion of all required actions in the Welcome panel. The standard Home topic will be hidden by the Welcome panel until all required actions are complete.

Using guided setup

The Welcome panel provides options for you to quickly and easily set up your system by guiding you through the firmware update and configuration process. With guided setup, you must first access the Update Firmware panel where you can review the controller module firmware version and perform recommended updates. When finished, you must configure your system settings by accessing the System Settings panel and completing all required options. Once these options are complete you can provision your system.

NOTE: A user with the `manage` role must complete the guided setup process.

The Welcome panel also displays the system's health. If the system's health is degraded or faulty, you can click System Information to access the System topic. Here you can view information about each enclosure, including its physical components, in front, rear, and tabular views. For more information on the System topic, see [“Working in the System topic” \(page 61\)](#).

If the system health is degraded, you may still be able to update your controller module firmware and configure the system. It is recommended that you resolve any health issues before updating firmware or configuring the system.

The Welcome panel displays until all required actions have been completed in the Update Firmware panel and the System Settings panel. Completed actions are marked with a check mark. You cannot access the Home topic until the Welcome panel actions are complete.

To use guided setup

1. As a user with the `manage` role, click **Upgrade Firmware** from the Welcome panel.
2. Review the firmware versions for the controller, enclosure, and drive firmware.
3. Per HPE recommendation, verify your firmware is up to date by clicking on the link to find the latest available firmware on the HPE support website. For more information, see [“Updating firmware” \(page 66\)](#).
4. When you are finished with the Update Firmware panel, click **Close**.
5. From the Welcome panel, click **System Settings**.
6. Choose options to configure your system. For more information, see [“Configuring system settings” \(page 41\)](#).

NOTE: Tabs with a red asterisk next to them are required.

7. Save your settings and exit System Settings to return to the Welcome panel.
8. Click **Go to Home topic** to access the Home topic.

Viewing overall system status

The Home topic provides an overview of the storage managed by the system. Information is shown about hosts, host ports, storage capacity and usage, global spares, and logical storage components (like volumes, snapshots, disk groups, and pools).

Host information










The Hosts block shows how many host groups, hosts, and initiators are defined in the system. An *initiator* identifies 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. A *host* is a user-defined set of initiators that represents a server. A *host group* is a user-defined set of hosts for ease of management.

NOTE: If the external port is a switch and there is no connection from the switch to an I/O adapter, then no host information will be shown.

Port information

The Ports A block shows the name and type (protocol) of each host port in controller A. The port icon indicates whether the port is active or inactive:

	FC port is active.
	FC port is connected.
	FC port is disconnected.
	iSCSI port is active.
	iSCSI port is connected.
	iSCSI port is disconnected.
	SAS port is active.
	SAS port is connected.
	SAS port is disconnected.

The Ports B block shows similar information for controller B.

Hover the cursor over a port to see the following information in the Port Information panel. If the health is not OK, the health reason and recommended action are shown to help you resolve problems.

Port Information	<p>FC port: Name, type, ID (WWN), status, configured speed, actual speed, topology, primary loop ID, supported speeds, SFP status, part number, and health</p> <p>iSCSI IPv4 port: Name, type, ID (IQN), status, configured speed, actual speed, IP version, MAC address, IP address, gateway, netmask, SFP status, part number, 10G compliance, cable length, cable technology, Ethernet compliance, and health</p> <p>SAS port: Name, type, ID (WWN), status, actual speed, topology, expected lanes, active lanes, disabled lanes, cable type, and health</p>
------------------	--

The area between the blocks displays the following statistics that show the current performance from all hosts to the system:

- Current IOPS for all ports, calculated over the interval since these statistics were last requested (every 30 seconds unless more than one SMU session is active or if the CLI command `show host-port-statistics` is issued) or reset.
- Current data throughput (MB/s) for all ports, calculated over the interval since these statistics were last requested or reset.

For MSA 1050: For a system with a 2-port SAS controller module, host ports can be configured to use fan-out cables or standard cables. If fan-out cables are connected to SAS ports that are configured to use them, fan-out cable icons appear between the depicted SAS ports.

Capacity information

The Capacity block shows two color-coded bars. The lower bar represents the physical capacity of the system, showing the capacity of disk groups, global spares, and unused disk space, if any. The upper bar identifies how the capacity is allocated and used. For color-code descriptions, see [“Color codes” \(page 14\)](#).

The upper bar shows the reserved, allocated, and unallocated space for the system. Reserved space refers to space that is unavailable for host use. It consists of RAID parity and the metadata needed for internal management of data structures. The terms allocated space and unallocated space have the following meanings:

- Allocated space is the amount of space that the data written to the pools takes.
- Unallocated space is space that is designated for a pool but has not yet been allocated by a volume within that pool.
- Uncommitted space is the overall space minus the allocated and unallocated space.

If virtual storage is *overcommitted*, which means that the amount of storage capacity that is designated for use by volumes exceeds the physical capacity of the storage system, then the right upper bar will be longer than the lower bar.

Hover the cursor over a segment of a bar to see the storage size represented by that segment. Point anywhere in this block to see the following information about capacity utilization in the Capacity Utilization panel:

- Total Disk Capacity. The total physical capacity of the system
- Unused. The total unused disk capacity of the system
- Global Spares. The total global spare capacity of the system
- Virtual Disk Groups. The capacity of disk groups, both total and by pool
- Reserved. The reserved space for disk groups, both total and by pool
- Allocated. The allocated space for disk groups, both total and by pool
- Unallocated. The unallocated space for disk groups, both total and by pool
- Uncommitted. The uncommitted space in each pool (total space minus the allocated and unallocated space) and total uncommitted space

Storage information

The Storage A and Storage B blocks provide more detailed information about the logical storage of the system. The Storage A block shows information about virtual pool A, which is owned by controller A. The Storage B block shows the same types of information about virtual pool B.

Each storage block contains color-coded graphs. For color-code descriptions, see [“Color codes” \(page 14\)](#).

For virtual storage, the block contains a pool capacity graph, a disk group utilization graph, and—if read cache is configured—a cache utilization graph. The pool capacity graph consists of two horizontal bars. The top bar represents the allocated and unallocated storage for the pool with the same information as the capacity top bar graph, but for the pool instead of the system. The bottom horizontal bar represents the size of the pool.

The disk group utilization graph consists of a graph with vertical measurements. The size of each disk group in the virtual pool is proportionally represented by a horizontal section of the graph. Vertical shading for each disk group section represents the relative space allocated in that disk group. A tool tip for each section shows the disk group name, size, and amount of unallocated space. The color for each disk group represents the tier to which it belongs.

The cache utilization graph also consists of a graph with vertical measurements. However, since read cache does not cache pool capacity, it is represented independently.

The number of volumes and snapshots for the pool owned by the controller displays above the top horizontal bar.

Hover the cursor anywhere in a storage block to display the Storage Information panel.

Storage Information for a virtual pool	Owner, storage type, total size, allocated size, snapshot size, available size, allocation rate, and deallocation rate For each tier: Pool percentage, number of disks, total size, allocated size, unallocated size, number of reclaimed pages, and health If the pool health is not OK, an explanation and recommendations for resolving problems with unhealthy components is available. If the overall storage health is not OK, the health reason, recommended action, and unhealthy subcomponents are shown to help you resolve problems.
--	---

System health information

The health icon between the storage blocks indicates the health of the system. Hover the cursor over this icon to display the System Health panel, which shows more information about the health state. If the system health is not OK, the System Health panel also shows information about resolving problems with unhealthy components.

Spares information

The Spares block between the storage blocks and below the event icon shows the number of disks that are designated as global spares to automatically replace a failed disk in the system. Hover the cursor over the Spares block to see the disk types of the available global spares in the Global Spares Information panel.

Resolving a pool conflict caused by inserting a foreign disk group

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

This is not a common operation, and you should consider your conflict resolution options carefully. To resolve this conflict, do either of the following:

- If the pool conflict was expected—for example, you want to access data from the disk group from `OldSystem`'s pool A—unmount and unmap the LUNs from any host accessing volumes on `NewSystem`, stop I/O from hosts accessing any volumes on `NewSystem`, and power down `NewSystem`. Then, physically remove all disks for `NewSystem`'s original pool A. Restore power to `NewSystem` (with disk groups from `OldSystem`'s pool A still inserted in `NewSystem`). The data from the disk groups from `OldSystem`'s pool A will be accessible; copy that data to pool B on `NewSystem`. After you have copied the data (volumes), you can safely swap the `OldSystem` disks back out of the system and replace them with the `NewSystem` disks. Remap and remount the LUNs to any host that requires access to volumes on `NewSystem`'s pool A.

⚠ CAUTION: This is an offline operation. Removing a virtual disk group or pool while the system is online may result in corruption and possible data loss. Power off the system before removing physical disks.

- If the pool conflict was unexpected—for example, you did not realize a previous pool existed on the `OldSystem` disks and data contained on the disks is no longer needed—remove the `OldSystem` disks out of `NewSystem`, put the disks back into `OldSystem`, delete the pool off the `OldSystem` disks, and then re-insert the disks into `NewSystem`. The `OldSystem` disks will now show as available and can be added to an existing pool on `NewSystem`.

⚠ CAUTION: Deleting a pool will delete all the data it contains.

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

Configuring system settings

The System Settings panel provides options for you to quickly and easily configure your system, including:

- Setting the system date and time
- Managing users
- Installing system licenses
- Configuring controller management network ports
- Enabling or disabling system-management services
- Entering system identification information
- Setting system notification settings
- Configuring host ports (if applicable)

Access the panel by doing one of the following:

- In the Home topic, select **Action > System Settings**.
- In the System topic, select **Action > System Settings**.
- In the Welcome panel, select **System Settings**.

Navigate the options by clicking the tabs located on the left side of the panel. Tabs with a red asterisk next to them are required. You can apply changes by clicking **Apply** at any point to save your settings. You can apply changes and close the panel by clicking **Apply and Close**.

 **TIP:** It is recommended that you click Apply before moving to a new tab in the System Settings panel.

Setting the system date and time

Use the Date and Time panel to change the storage system date and time that appear in the banner. It is important to set the date and time so that entries in system logs and notifications have correct time stamps.

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

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

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

Whether NTP is enabled or disabled, the storage system does not automatically make time adjustments for Daylight Saving Time. You must make such adjustments manually.

To manually enter date and time settings

1. Perform one of the following to access the Date and Time options:
 - In the Home topic, select **Action > System Settings > Date and Time**.
 - In the System topic, select **Action > System Settings > Date and Time**.
 - In the banner, click the System Date/Time Bar panel and select **Set Date and Time**.
 - In the Welcome panel, select **System Settings > Date and Time**.
2. If checked, clear the **Network Time Protocol (NTP)** check box.

3. To set the **Date** value, enter the current date in the format *YYYY-MM-DD*.
4. To set the **Time** value, enter two-digit values for the hour and minutes and select either **AM**, **PM**, or **24H** (24-hour clock).
5. Perform one of the following:
 - o To save your settings and continue configuring your system, click **Apply**.
 - o To save your settings and close the panel, click **Apply and Close**.A confirmation panel displays.
6. Click **OK** to save your changes. Otherwise, click **Cancel**.

To obtain the date and time from an NTP server

1. Perform one of the following to access the Date and Time options:
 - o In the Home topic, select **Action > System Settings > Date and Time**.
 - o In the System topic, select **Action > System Settings > Date and Time**.
 - o In the banner, click the System Date/Time Bar panel and select **Set Date and Time**.
 - o In the Welcome panel, select **System Settings > Date and Time**.
2. Select the **Network Time Protocol (NTP)** check box.
3. Perform one of the following:
 - o To have the system retrieve time values from a specific NTP server, enter its IP address in the NTP Server Address field.
 - o To have the system listen for time messages sent by an NTP server in broadcast mode, clear the NTP Server Address field.
4. In the NTP Time Zone Offset field, enter the time zone as an offset in hours, and optionally minutes, from UTC. For example: the Pacific Time Zone offset is -8 during Pacific Standard Time or -7 during Pacific Daylight Time and the offset for Bangalore, India is +5:30.
5. Perform one of the following:
 - o To save your settings and continue configuring your system, click **Apply**.
 - o To save your settings and close the panel, click **Apply and Close**.A confirmation panel displays.
6. Click **OK** to save your changes. Otherwise, click **Cancel**.

Managing users

The Manage Users panel lets you manage both local users and LDAP user groups. From here, you can add, modify, and delete users, determine user permissions, and set system preferences based on individual user profiles. Settings with a red asterisk next to them are required.

Local users

The system provides three default users and nine additional users can be created. The default users are “standard users,” which can access one or more of the following management interfaces: SMU, CLI, SMI-S, or FTP and SFTP. You can also create SNMPv3 users, which can either access the Management Information Base (MIB) or receive trap notifications. SNMPv3 users support SNMPv3 security features, such as authentication and encryption. For information about configuring trap notifications, see [“Setting system notification settings” \(page 52\)](#). For information about the MIB, see [“SNMP reference” \(page 150\)](#).

As a user with the manage role you can modify any user, or you can delete any user other than the current user. As a user with only a standard or monitor role, you can change settings for the current user with the exception of user interfaces and role. You also cannot change the settings of other user groups. The following table shows the settings for default users.

Table 9 Settings for the default users

User Name	Password	User Type	Roles	Interfaces	Base	Precision	Unit	Temperature	Timeout (minutes)	Locale
monitor	!monitor	Standard	monitor	WBI, CLI	Base 10	1	Auto	Celsius	30	English
manage	!manage		manage, standard, monitor	WBI, CLI, SMI-S, FTP, SFTP						
ftp	!ftp		manage, standard, monitor	FTP, SFTP						

! **IMPORTANT:** To secure the storage system, set a new password for each default user.

The following options apply to standard and SNMPv3 users:

- **User Name.** A user name 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: " , < \ :
- **Password.** A password is case sensitive and can have 8–32 characters. If the password contains only printable ASCII characters, then it must contain at least one uppercase character, one lowercase character, one numeric character, and one non-alphanumeric character. A password can include printable UTF-8 characters except for the following: a space or " ' , < > \
- **Confirm Password.** Re-enter the new password.
- **User Type.** When creating a new user, select **Standard** to show options for a standard user, or **SNMPv3** to show options for an SNMPv3 user. The default is Standard.

The following options apply only to a standard user:

- **Roles.** Select one or more of the following roles:
 - **Manage.** Enables the user to change system settings.
 - **Standard.** Enables the user to change system settings except for: creating or deleting local users, modifying user role and interfaces, configuring LDAP, performing write operations through FTP or SFTP, performing file uploads from the SMU (WBI), or using the CLI `restore defaults` command.
 - **Monitor.** Enables the user to view but not change system status and settings. This is enabled by default and cannot be disabled.
- **Interfaces.** Select one or more of the following interfaces:
 - **WBI.** Enables access to the SMU. This is a default.
 - **CLI.** Enables access to the command-line interface. This is a default.
 - **SMI-S.** Enables access to the SMI-S interface, which is used for remote management of the system through your network.
 - **FTP.** Enables access to the FTP interface or the SFTP interface, which can be used instead of the WBI to install firmware updates and to download logs.
- **Base Preference.** Select the base for entry and display of storage-space sizes:
 - **Base 2.** Sizes are shown as powers of 2, using 1024 as a divisor for each magnitude.
 - **Base 10.** Sizes are shown as powers of 10, using 1000 as a divisor for each magnitude. This is a default.
- **Precision Preference.** Select the number of decimal places (1–10) for display of storage-space sizes. The default is 1.

- **Unit Preference.** Select one of the following options for display of storage-space sizes:
 - **Auto.** Enables the system to determine the proper unit for a size. Based on the precision setting, if the selected unit is too large to meaningfully display a size, the system uses a smaller unit for that size. For example, if the unit is set to TB and the precision is set to 1, the size 0.11709 TB is shown as 117.1 GB. This is the default.
 - **TB.** Display all sizes in terabytes.
 - **GB.** Display all sizes in gigabytes.
 - **MB.** Display all sizes in megabytes.
- **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 (2–720 minutes). The default is 30 minutes.
- **Locale.** 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” \(page 15\)](#).

The following options apply only to an SNMPv3 user:

- **SNMPv3 Account Type.** Select one of the following types:
 - **User Access.** Enables the user to view the SNMP MIB. This is the default.
 - **Trap Target.** Enables the user to receive SNMP trap notifications.
- **SNMPv3 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.
- **SNMPv3 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.
- **SNMPv3 Privacy Password.** If the privacy type is set to use encryption, specify an encryption password. This password is case sensitive and can have 8–32 characters. If the password contains only printable ASCII characters, then it must contain at least one uppercase character, one lowercase character, and one non-alphabetic character. A password can include printable UTF-8 characters except for the following: a space or " , < > \
- **Trap Host Address.** If the account type is **Trap Target**, specify the network address of the host system that will receive SNMP traps. The value can be an IPv4 address.

Adding, modifying, and deleting users

To add a new user

1. As a user with a `manage` role, perform one of the following:
 - In the Home topic, select **Action > System Settings > Manage Users**.
 - In the System topic, select **Action > System Settings > Manage Users**.
 - In the banner, click the User panel and select **Manage Users**.
 - In the Welcome panel, select **System Settings > Manage Users**.

The Local Users tab displays a table of existing users and options to set. For information about using tables, see [“Tips for using tables” \(page 12\)](#).

2. Below the table, click **New**.
3. Set the options.
4. Perform one of the following:
 - To save your settings and continue configuring your system, click **Apply**.
 - To save your settings and close the panel, click **Apply and Close**.

A confirmation panel displays.

5. Click **OK** to save your changes. Otherwise, click **Cancel**.

To create a user from an existing user

1. As a user with a `manage` role, perform one of the following:
 - o In the Home topic, select **Action > System Settings > Manage Users**.
 - o In the System topic, select **Action > System Settings > Manage Users**.
 - o In the banner, click the User panel and select **Manage Users**.
 - o In the Welcome panel, select **System Settings > Manage Users**.

The Local Users tab displays a table of existing users and options to set. For information about using tables, see [“Tips for using tables” \(page 12\)](#).

2. Select the user to copy.
3. Click **Copy**. A user named `copy_of_selected-user` appears in the table.
4. Set a new user name and password and optionally change other settings.
5. Perform one of the following:
 - o To save your settings and continue configuring your system, click **Apply**.
 - o To save your settings and close the panel, click **Apply and Close**.A confirmation panel displays.
6. Click **OK** to save your changes. Otherwise, click **Cancel**.

To modify a user

1. As a user with a `manage` role, perform one of the following:
 - o In the Home topic, select **Action > System Settings > Manage Users**.
 - o In the System topic, select **Action > System Settings > Manage Users**.
 - o In the banner, click the User panel and select **Manage Users**.
 - o In the Welcome panel, select **System Settings > Manage Users**.

The Local Users tab displays a table of existing users and options to set. For information about using tables, see [“Tips for using tables” \(page 12\)](#).

NOTE: Users with a `standard` or `monitor` role can modify their own settings and preferences using the CLI commands `set user` and `set user-group`.

2. Select the user to modify.
3. Change the settings. You cannot change the user name.
4. Perform one of the following:
 - o To save your settings and continue configuring your system, click **Apply**.
 - o To save your settings and close the panel, click **Apply and Close**.A confirmation panel displays.
5. Click **OK** to save your changes. Otherwise, click **Cancel**.

To delete a user (other than your current user)

1. As a user with a `manage` role, perform one of the following:
 - o In the Home topic, select **Action > System Settings > Manage Users**.
 - o In the System topic, select **Action > System Settings > Manage Users**.
 - o In the banner, click the User panel and select **Manage Users**.
 - o In the Welcome panel, select **System Settings > Manage Users**.

The Local Users tab displays a table of existing users and options to set. For information about using tables, see [“Tips for using tables” \(page 12\)](#).

2. Select the user to delete.

3. Click **Delete**. A confirmation panel appears.
4. Click **OK** to delete the user. Otherwise, click **Cancel**. If you clicked OK, the user is removed, the table is updated, and any sessions associated with that user name are terminated.

NOTE: The system requires at least one user with the `manage` role to exist.

LDAP users

The LDAP Users tab lets you enable and disable LDAP configuration. When LDAP is enabled, LDAP users with a `manage` role can create, modify, and delete both local users and user groups. Local users with a `manage` role will be able to create and delete user groups whether or not LDAP is enabled. LDAP users can access one or more of the following management interfaces: the WBI, CLI, FTP or SMI-S.

Users with a `manage` role can 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 the storage size base, precision, and units; temperature-scale; timeout; and locale.

Users logging into the WBI 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/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 Configuration checkbox and enter the Server address, Port, and User-search-base. If the Port is left blank it will default to 636. For more information about these options, see [“About LDAP” \(page 31\)](#).

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 type, role, and interfaces. You also cannot change the settings of other user groups.

Panel options are defined as follows:

- Roles. Select one or more of the following roles:
 - **Manage**. Enables the user group to change system settings.
 - **Standard**. Enables the user group to change system settings except for: creating, modifying, or deleting user or user groups, performing write operations through FTP or SFTP, performing file uploads from the SMU (WBI), or using the CLI `restore defaults` command.
 - **Monitor**. Enables the user group to view but not change system status and settings. This is enabled by default and cannot be disabled.
- Interfaces. Select one or more of the following interfaces:
 - **WBI**. Enables access to the WBI. This is a default.
 - **CLI**. Enables access to the command-line interface. This is a default.
 - **FTP**. Enables access to the FTP interface or the SFTP interface, which can be used instead of the WBI to install firmware updates and to download logs.
 - **SMI-S**. Enables access to the SMI-S interface, which is used for remote management of the system through your network.
- Base Preference. Select the base for entry and display of storage-space sizes:
 - **Base 2**. Sizes are shown as powers of 2, using 1024 as a divisor for each magnitude.
 - **Base 10**. Sizes are shown as powers of 10, using 1000 as a divisor for each magnitude. This is a default.
- Precision Preference. Select the number of decimal places (1–10) for display of storage-space sizes. The default is 1.

- **Unit Preference.** Select one of the following options for display of storage-space sizes:
 - **Auto.** Enables the system to determine the proper unit for a size. Based on the precision setting, if the selected unit is too large to meaningfully display a size, the system uses a smaller unit for that size. For example, if the unit is set to TB and the precision is set to 1, the size 0.11709 TB is shown as 117.1 GB. This is the default.
 - **TB.** Display all sizes in terabytes.
 - **GB.** Display all sizes in gigabytes.
 - **MB.** Display all sizes in megabytes.
- **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 (2–720 minutes). The default is 30 minutes.
- **Locale.** 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” \(page 15\)](#).

To enable LDAP

1. As a user with a `manage` role, perform one of the following:
 - In the Home topic, select **Action > System Settings > Manage Users**, then click the **LDAP Users** tab.
 - In the System topic, select **Action > System Settings > Manage Users**, then click the **LDAP Users** tab.
 - In the banner, click the User panel and select **Manage Users**, then click the **LDAP Users** tab.
 - In the Welcome panel, select **System Settings > Manage Users**, then click the **LDAP Users** tab.
 The LDAP Users tab displays configuration parameters to set and a table of existing user groups.
2. Click **Enable LDAP Configuration**.
3. Enter the **Server** address, **Port**, and **User-search-base**. For more information about these options, see [“About LDAP” \(page 31\)](#).

NOTE: If you do not enter a Port, the system will use 636 as the default.

4. Optional: Enter the **Alt-Server** address and the **Alt-Port** address.
5. Perform one of the following:
 - To save your settings and continue configuring your system, click **Apply**.
 - To save your settings and close the panel, click **Apply and Close**.
 A confirmation panel displays.
6. Click **OK** to save your changes. Otherwise, click **Cancel**.

To create a new LDAP user group

1. As a user with a `manage` role, perform one of the following:
 - In the Home topic, select **Action > System Settings > Manage Users**, then click the **LDAP Users** tab.
 - In the System topic, select **Action > System Settings > Manage Users**, then click the **LDAP Users** tab.
 - In the banner, click the User panel and select **Manage Users**, then click the **LDAP Users** tab.
 - In the Welcome panel, select **System Settings > Manage Users**, then click the **LDAP Users** tab.
 The LDAP Users tab displays configuration parameters to set and a table of existing user groups.
2. Below the Current User-Groups table, click **New**.
3. Enter the name of the new user group, then set the options.

4. Perform one of the following:
 - o To save your settings and continue configuring your system, click **Apply**.
 - o To save your settings and close the panel, click **Apply and Close**.

A confirmation panel displays.

5. Click **OK** to save your changes. Otherwise, click **Cancel**.

To modify a LDAP user group

1. As a user with a `manage` role, perform one of the following:
 - o In the Home topic, select **Action > System Settings > Manage Users**, then click the **LDAP Users** tab.
 - o In the System topic, select **Action > System Settings > Manage Users**, then click the **LDAP Users** tab.
 - o In the banner, click the User panel and select **Manage Users**, then click the **LDAP Users** tab.
 - o In the Welcome panel, select **System Settings > Manage Users**, then click the **LDAP Users** tab.

The LDAP Users tab displays configuration parameters to set and a table of existing user groups.

2. Select the user group to modify.
3. Change the settings.
4. Perform one of the following:
 - o To save your settings and continue configuring your system, click **Apply**.
 - o To save your settings and close the panel, click **Apply and Close**.

A confirmation panel displays.

5. Click **OK** to save your changes. Otherwise, click **Cancel**.

To delete a LDAP user group

1. As a user with a `manage` role, perform one of the following:
 - o In the Home topic, select **Action > System Settings > Manage Users**, then click the **LDAP Users** tab.
 - o In the System topic, select **Action > System Settings > Manage Users**, then click the **LDAP Users** tab.
 - o In the banner, click the User panel and select **Manage Users**, then click the **LDAP Users** tab.
 - o In the Welcome panel, select **System Settings > Manage Users**, then click the **LDAP Users** tab.

The LDAP Users panel displays configuration parameters to set and a table of existing user groups.

2. Select the user group to delete.

NOTE: You cannot delete a user group that is logged into the system.

3. Click **Delete**. A confirmation panel appears.
4. Click **OK** to continue. Otherwise, click **Cancel**. If you clicked **OK**, the user group is removed and the table is updated.

Installing a license

A license is required to use the Performance tier, expand the maximum number of snapshots, and use the replication feature. The license is specific to a controller enclosure and firmware version.

NOTE: The system supports use of only spinning disks or only SSDs. A Performance Tier license is required to use a combination of both. For information about all-flash array feature, see [“All-flash array” \(page 20\)](#). For information about the rules for using SSDs and spinning disks, see [“About SSDs” \(page 20\)](#).

A base license is permanently installed to use virtualization, up to 64 snapshots, volume copy, and Volume Shadow Copy Service (VSS).

Viewing the status of licensed features

1. Perform one of the following to access the Install License options:
 - o In the Home topic, select **Action > System Settings > Install License**.
 - o In the System topic, select **Action > System Settings > Install License**.
 - o In the Welcome panel, select **System Settings > Install License**.
2. View the status of the License Settings:
 - o Feature. The feature name.
 - o Base. One of the following:
 - The number of standard snapshots that users can create without a license.
 - N/A. Not applicable.
 - o License. One of the following:
 - The number of standard snapshots that the installed license supports.
 - Enabled. The feature is enabled.
 - Disabled. The feature is disabled.
 - o In Use. One of the following:
 - The number of standard snapshots that exist.
 - N/A. Not applicable.
 - o Max Licensable. One of the following:
 - The number of standard snapshots that the maximum license supports.
 - N/A. Not applicable.
 - o Expiration. One of the following:
 - Never. License does not expire.
 - N/A. Not applicable.

The panel also shows the licensing serial number and the licensing version number (both required for generating a license).

Installing a license

1. Verify the following:
 - o The license file is saved to a network location that you can access from the SMU.
 - o You are signed into the controller enclosure for which the file is generated.
2. Perform one of the following to access the Install License options:
 - o In the Home topic, select **Action > System Settings > Install License**.
 - o In the System topic, select **Action > System Settings > Install License**.
 - o In the Welcome panel, select **System Settings > Install License**.
3. Click **Choose File** to locate and select the license file.
4. Perform one of the following:
 - o To save your settings and continue configuring your system, click **Apply**.
 - o To save your settings and close the panel, click **Apply and Close**.A confirmation panel displays.
5. Click **OK** to save your changes. Otherwise, click **Cancel**. If you clicked OK, the license settings table is updated and, for each feature included in the license, the *Expiration* value changes to *Never*.

Configuring controller network ports

You can manually set static IP addressing parameters for network ports or you can specify that IP values be set automatically using DHCP for IPv4. When setting IP values, you can choose IPv4 formatting for each controller.

When using DHCP mode, the system obtains values for the network port IP address, subnet mask, and gateway from a DHCP server if one is available. If a DHCP server is unavailable, the system will use its default values. You must have some means of determining what addresses have been assigned, such as the list of bindings on the DHCP server. You can retrieve the DHCP assigned IP addresses either through the USB serial connection using CLI commands or from the DHCP server list of MAC address to IP address leases.

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: 10.0.0.1

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

To set IPv4 values for network ports

1. Perform one of the following to access Network options:
 - In the Home topic, select **Action > System Settings**, then click the **Network** tab.
 - In the System topic, select **Action > System Settings**, then click the **Network** tab.
 - In the Welcome panel, select **System Settings**, and then click the **Network** tab.
2. Select the type of IP address to use for each controller:
 - Choose **Addressing Mode > manual** to enter static IP addresses.
 - Choose **Addressing Mode > DHCP** to have the system automatically obtain values from a DHCP server.
3. If you chose manual, enter the unique IP address, IP mask, and gateway values for each controller, then record the IP values you assign.

NOTE: The following IP addresses are reserved for internal use by the storage system: 169.254.255.1, 169.254.255.2, 169.254.255.3, 169.254.255.4, and 127.0.0.1. Because these addresses are routable, do not use them anywhere in your network.

4. Perform one of the following:
 - To save your settings and continue configuring your system, click **Apply**.
 - To save your settings and close the panel, click **Apply and Close**.A confirmation panel displays.
5. Click **OK** to continue. If you chose DHCP and the controllers successfully obtained IP values from the DHCP server, the new IP values appear. Record the new addresses and sign out to use the new IP address to access the WBI.

Enabling or disabling system-management services

You can enable or disable management 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. In-band services operate through the data path and can slightly reduce I/O performance. To allow specific users to access the SMU, CLI, or other interfaces, see [“Adding, modifying, and deleting users” \(page 44\)](#).

To change system services settings

1. Perform one of the following to access Services options:
 - o In the Home topic, select **Action > System Settings**, then click the **Services** tab.
 - o In the System topic, select **Action > System Settings**, then click the **Services** tab.
 - o In the banner, click the System panel and select **Set Up System Services**.
 - o In the Welcome panel, select **System Settings**, and then click the **Services** tab.
2. Enable the services that you want to use to manage the storage system, and disable the others.
 - o Web Browser Interface (WBI). The web application that is the primary interface for managing the system.
 - o You can enable use of HTTP, HTTPS for increased security, or both. If you disable both, you will lose access to the WBI interface. By default, HTTP is disabled while HTTPS is enabled.
 - o Command Line Interface (CLI). An advanced-user interface that is used to manage the system and can be used to write scripts. You can enable use of SSH (secure shell) for increased security, Telnet, or both. If you select SSH, specify the port number to use. The default is 22. By default, Telnet is disabled and SSH is enabled.
 - o Storage Management Initiative Specification (SMI-S). Used for remote management of the system through your network. You can enable use of secure (encrypted) or unsecure (unencrypted) SMI-S:
 - **Enable**. Select this check box to enable unencrypted communication between SMI-S clients and the embedded SMI-S provider in each controller module via HTTP port 5988. Clear this check box to disable the active port and use of SMI-S.
 - **Encrypted**. Select this check box to enable encrypted communication, which disables HTTP port 5988 and enables HTTPS port 5989 instead. Clear this check box to disable port 5989 and enable port 5988. This is the default.
 - o **Service Location Protocol (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 enabled by default.
 - o **File Transfer Protocol (FTP)**. A secondary interface for installing firmware updates and downloading logs. FTP is disabled by default.
 - o **SSH File Transfer Protocol (SFTP)**. A secure secondary interface for installing firmware updates, downloading logs, and installing security certificates and keys, All data sent between the client and server will be encrypted. SFTP is enabled by default. If selected, specify the port number to use. The default is 1022.
 - o **Simple Network Management Protocol (SNMP)**. Used for remote monitoring of the system through your network.
 - o **Service Debug**. Used for technical support only. Enables or disables debug capabilities, including Telnet debug ports and privileged diagnostic user IDs. This is disabled by default. Enabling the service debug interface allows remote connection, through incoming ports only, by HPE or HPE's authorized representatives for troubleshooting. Disabling the service debug interface removes this access.
 - o **Activity Progress Reporting**. Provides access to the activity progress interface via HTTP port 8081. This mechanism reports whether a firmware update or partner firmware update operation is active and shows the progress through each step of the operation. In addition, when the update operation completes, status is presented indicating either the successful completion, or an error indication if the operation failed.
 - o **In-band SES Capability**. Used for in-band monitoring of system status based on SCSI Enclosure Services (SES) data. This service operates through the data path and can slightly reduce I/O performance. SES is disabled by default.
3. Perform one of the following:
 - o To save your settings and continue configuring your system, click **Apply**.
 - o To save your settings and close the panel, click **Apply and Close**.A confirmation panel displays.
4. Click **OK** to save your changes. Otherwise, click **Cancel**.


Entering system identification information

To change system information settings

1. Perform one of the following to access the Services options:
 - o In the Home topic, select **Action > System Settings**, then click the **System information** tab.
 - o In the System topic, select **Action > System Settings**, then click the **System information** tab.
 - o In the banner, click the System panel and select **Set System Information**.
 - o In the Welcome panel, select **System Settings**, and then click the **System information** tab.
2. Set the system name, contact, location, and information (description) values. The name is shown in the browser title bar or tab. The name, location, and contact are included in event notifications. All four values are recorded in system debug logs for reference by service personnel. Each value can include a maximum of 79 bytes, using all characters except the following: " < > \
A confirmation panel displays.
3. Perform one of the following:
 - o To save your settings and continue configuring your system, click **Apply**.
 - o To save your settings and close the panel, click **Apply and Close**.
4. Click **OK** to save your changes. Otherwise, click **Cancel**.

Setting system notification settings

The Notifications tab provides options for you to set up and test several types of system notifications. These include:

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

- Configuring SMTP settings.
- Sending notifications to email addresses when events occur in the system. The system can also be configured to send weekly alerts about system health issues to configured email addresses until corrective action has been taken and the system health value has returned to OK.
- Sending notifications to SNMP trap hosts.
- Enabling managed logs settings, which transfers log data to a log-collection system. For more information about the managed logs feature, see [“About managed logs” \(page 30\)](#).
- Setting remote syslog notifications to allow events to be logged by the syslog of a specified host computer. Syslog is 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).
- Testing notifications.

NOTE: Settings with a red asterisk next to them are required.

To configure SMTP settings

1. Perform one of the following to access the options in the Notifications tab:
 - o In the Home topic, select **Action > System Settings**, then click the **Notifications** tab.
 - o In the System topic, select **Action > System Settings**, then click the **Notifications** tab.
 - o In the footer, click the events panel and select **Set Up Notifications**.
 - o In the Welcome panel, select **System Settings**, and then click the **Notifications** tab.
2. If the mail server is not on the local network, make sure that the gateway IP address was set in [“Configuring controller network ports” \(page 49\)](#).
3. Select the **Email** tab.

4. In the SMTP Server address field, enter the IP address of the SMTP mail server to use for the email messages.
5. In the Sender Domain field, enter a domain name, which will be joined with an @ symbol to the sender name to form the “from” address for remote notification. The domain name can have a maximum of 255 bytes. Because this name is used as part of an email address, do not include spaces or the following: \ " ; < > ()
The default is mydomain.com. If the domain name is not valid, some email servers will not process the mail.
6. In the Sender Name field, enter a sender name, which will be joined with an @ symbol to the domain name to form the “from” address for remote notification. This name provides a way to identify the system that is sending the notification. The sender name can have a maximum of 64 bytes. Because this name is used as part of an email address, do not include spaces or the following: \ " ; < > () []
For example: Storage-1.
7. In the Port text box, enter the port to use when communicating with the SMTP server. Leaving the port field blank tells the system to use the default port associated with the security protocol selected in the following step.
8. Set the security protocol to use when communicating with the SMTP server:
 - o **None.** Does not use a security protocol. The standard SMTP port is 25, and is the system default.
 - o **TLS.** Enables Transport Layer Security (TLS) authentication. The standard ports are 25 or 587. The system default is 587.
 - o **SSL.** Enables Secure Sockets Layer (SSL) authentication. The standard port is 465, the system default.
9. If you selected TLS or SSL, enter the password of the user in the Sender Name field, then confirm the password.
10. Perform one of the following:
 - o To save your settings and continue configuring your system, click **Apply**.
 - o To save your settings and close the panel, click **Apply and Close**.
 A confirmation panel displays.
11. Click **OK** to save your changes. Otherwise, click **Cancel**.

To send email notifications

1. Perform one of the following to access the options in the Notifications tab:
 - o In the Home topic, select **Action > System Settings**, then click the **Notifications** tab.
 - o In the System topic, select **Action > System Settings**, then click the **Notifications** tab.
 - o In the footer, click the events panel and select **Set Up Notifications**.
 - o In the Welcome panel, select **System Settings**, and then click the **Notifications** tab.
2. Select the **Email** tab and ensure that the SMTP Server and SMTP Domain options are set, as described in [“To configure SMTP settings” \(page 52\)](#).
3. Set the email notification:
 - o To enable email notifications, select the **Enable Email Notifications** check box. This enables the notification level and email address fields.
 - o To disable email notifications, clear the **Enable Email Notifications** check box. This disables the notification level and email address fields. This is the default.
4. If email notification is enabled, select the minimum severity for which the system should send email notifications: **Critical** (only); **Error** (and Critical); **Warning** (and Error and Critical); **Resolved** (and Error, Critical, and Warning); **Informational** (all).
5. If email notification is enabled, in one or more of the Email Address fields enter an email address to which the system should send notifications. Each email address must use the format user-name@domain-name. Each email address can have a maximum of 320 bytes. For example: Admin@mydomain.com or IT-team@mydomain.com.
6. If email notification is enabled:
 - o To enable sending weekly alerts about system health issues to configured email addresses on Sunday at 12:01 AM, select the **Enable Health Alerts** check box. This is the default.
 - o To disable sending weekly alerts about system health issues to configured email addresses, clear the **Enable Health Alerts** check box.

7. Perform one of the following:
 - o To save your settings and continue configuring your system, click **Apply**.
 - o To save your settings and close the panel, click **Apply and Close**.

A confirmation panel displays.

8. Click **OK** to save your changes. Otherwise, click **Cancel**.

To send notifications to SNMP trap hosts

1. Perform one of the following to access the options in the Notifications tab:
 - o In the Home topic, select **Action > System Settings**, then click the **Notifications** tab.
 - o In the System topic, select **Action > System Settings**, then click the **Notifications** tab.
 - o In the footer, click the events panel and select **Set Up Notifications**.
 - o In the Welcome panel, select **System Settings**, and then click the **Notifications** tab.
2. Select the **SNMP** tab. If a message near the top of the panel informs you that the SNMP service is disabled, enable it as shown in [“Enabling or disabling system-management services” \(page 50\)](#).
3. Select the minimum Notification Level severity for which the system should send email notifications: **Critical** (only); **Error** (and Critical); **Warning** (and Error and Critical); **Informational/Resolved** (all); or **none** (disabled, this is the default).
4. In the Read community field, enter the SNMP read password for your network. This password is included in traps that are sent. The value is case sensitive and can have a maximum of 31 bytes. It can include any character except for the following: " < >
The default is public.
5. In the Write community field, enter the SNMP write password for your network. The value is case sensitive and can have a maximum of 31 bytes. It can include any character except for the following: " ' < >
The default is private.
6. In the Trap Host Address fields enter the IP addresses of hosts that are configured to receive SNMP traps. The default is 0.0.0.0.
7. Perform one of the following:
 - o To save your settings and continue configuring your system, click **Apply**.
 - o To save your settings and close the panel, click **Apply and Close**.

A confirmation panel displays.

8. Click **OK** to save your changes. Otherwise, click **Cancel**.

To enable managed logs settings

1. Perform one of the following to access the options in the Notifications tab:
 - o In the Home topic, select **Action > System Settings**, then click the **Notifications** tab.
 - o In the System topic, select **Action > System Settings**, then click the **Notifications** tab.
 - o In the footer, click the events panel and select **Set Up Notifications**.
 - o In the Welcome panel, select **System Settings**, and then click the **Notifications** tab.
2. Select the **Email** tab and ensure that the SMTP Server and SMTP Domain options are set, as described in [“To configure SMTP settings” \(page 52\)](#).
3. Select the **Managed Logs** tab.
4. Set the managed log option:
 - o To enable managed logs, select the **Enable Managed Logs** check box.
 - o To disable managed logs, clear the **Enable Managed Logs** check box. This is the default.
5. If the managed logs option is enabled, in the Email destination address field, enter the email address of the log-collection system. The email address must use the format `user-name@domain-name` and can have a maximum of 320 bytes. For example: `LogCollector@mydomain.com`.

6. Select one of the following options:
 - o To use push mode, which automatically attaches system log files to managed-logs email notifications that are sent to the log-collection system, select the **Include logs as an email attachment** check box.
 - o To use pull mode, clear the **Include logs as an email attachment** check box. This is the default.
7. Perform one of the following:
 - o To save your settings and continue configuring your system, click **Apply**.
 - o To save your settings and close the panel, click **Apply and Close**.A confirmation panel displays.
8. Click **OK** to save your changes. Otherwise, click **Cancel**.

To set remote syslog notifications

1. Perform one of the following to access the options in the Notifications tab:
 - o In the Home topic, select **Action > System Settings**, then click the **Notifications** tab.
 - o In the System topic, select **Action > System Settings**, then click the **Notifications** tab.
 - o In the footer, click the events panel and select **Set Up Notifications**.
 - o In the Welcome panel, select **System Settings**, and then click the **Notifications** tab.
2. Select the **Syslog** tab.
3. Set the Syslog options:
 - o **Notification Level**. Select the minimum severity for which the system should send notifications: **Critical** (only); **Error** (and Critical); **Warning** (and Error and Critical); **Resolved** (and Error, Critical, and Warning); **Informational** (all); or **none** (Disabled), which disables syslog notification.
 - o **Syslog Server IP Address**. IP address of the syslog host system.
 - o **Syslog Server Port Number**. Port number of the syslog host system.
4. Perform one of the following:
 - o To save your settings and continue configuring your system, click **Apply**.
 - o To save your settings and close the panel, click **Apply and Close**.A confirmation panel displays.
5. Click **OK** to save your changes. Otherwise, click **Cancel**.

To test notification settings

1. For the purpose of receiving test notifications, which are sent with Informational severity, set the notification level to Informational as described in the procedures above.
2. Send test notifications to validate your selections by doing the following:
 - o On the Email tab, click **Send Email**. A test notification is sent to each email address.
 - o On the SNMP tab, click **Send SNMP**. A test notification is sent to each configured trap host.
 - o On the Managed Logs tab, click **Test Managed Logs**. A test notification is sent to the log-collection system.
 - o On the Syslog tab, click **Test Syslog**. A test notification is sent to the syslog server.A confirmation displays.
3. Verify that the test notification reached the intended location.
4. Click **OK** to confirm.

NOTE: If there was an error in sending a test notification, event 611 displays in the confirmation. For more information see the Event Guide.

5. Configure your system to receive notifications based on your preferred severity as described in the procedures above.

Changing host port settings

You can configure controller host-interface settings for ports except for systems with a 4-port SAS controller module. To enable the system to communicate with hosts, you must configure the system's host-interface options.

NOTE: If the current settings are correct, port configuration is optional.

For a system with a 2-port SAS controller module, host ports can be configured to use fan-out cables or standard cables. A fan-out cable can connect one port on each of two SAS hosts to one controller port, using two dedicated PHY lanes per port. A standard cable can connect one port on a SAS host to one controller port, using four PHY lanes per port. Use of fan-out cables is enabled by default. When configuring the host-interface settings for a 2-port SAS controller module, the Host Ports Settings panel displays the current link speed, cable type, number of PHY lanes expected for the SAS port, and number of PHY lanes active for each SAS port. The number of ports that display depends on the configuration.

❗ **IMPORTANT:** Changing the fan-out setting will change the logical numbering of controller host ports, which will cause port IDs in mappings between volumes and initiators to be incorrect. Therefore, before changing the fan-out setting, unmap all mappings. After you have changed the fan-out setting and connected the appropriate cables, you can re-create the mappings.

SAN host ports can be configured as all FC or all iSCSI ports, or a combination of both. FC ports support use of qualified 8-Gbit/s or 16-Gbit/s SFPs. You can set FC ports to auto-negotiate the link speed or to use a specific link speed. iSCSI ports support use of qualified 1-Gbit/s, 10-Gbit/s SFPs or qualified 10-Gbit/s Direct Attach Copper (DAC) cables. iSCSI port speeds are auto-negotiated.

Host ports for MSA 1050 are configured to 8-Gbit/s FC, 10-Gbit/s iSCSI, or 1-Gbit/s iSCSI with the appropriate, qualified SFPs installed. The protocol (FC or iSCSI) and speed cannot be changed. 10-Gbit/s SFPs may be removed in favor of qualified Direct Attached Copper (DAC) cables.

For more information on MSA 2050 supported host port configurations, see www.hpe.com/support/msa2050QuickSpecs.

For more information on MSA 1050 supported host port configurations, see www.hpe.com/support/msa1050QuickSpecs.

NOTE: For information about setting host parameters such as FC port topology, and the host-port mode, see the CLI Reference Guide.

To configure FC ports

1. Perform one of the following to access the options in the Ports tab:
 - o In the Home topic, select **Action > System Settings**, then click the **Ports** tab.
 - o In the System topic, select **Action > System Settings**, then click the **Ports** tab.
 - o In the Welcome panel, select **System Settings**, and then click the **Ports** tab.
2. From the Host Post Mode dropdown, select FC.
3. From the Port Settings tab, set the port-specific options:
 - o Set the Speed option to the proper value to communicate with the host, or to auto, which auto-negotiates the proper link speed. Because a speed mismatch prevents communication between the port and host, set a speed only if you need to force the port to use a known speed. The maximum link speed is determined by the installed SFP.
 - o Set The FC Connection Mode to either point-to-point or auto:
 - **point-to-point:** Fibre Channel point-to-point. This is the default.
 - **auto:** Automatically sets the mode based on the detected connection type.

4. Perform one of the following:
 - o To save your settings and continue configuring your system, click **Apply**.
 - o To save your settings and close the panel, click **Apply and Close**.

A confirmation panel displays.

5. Click **OK** to save your changes. Otherwise, click **Cancel**.

To configure iSCSI ports

1. Perform one of the following to access the options in the Ports tab:
 - o In the Home topic, select **Action > System Settings**, then click the **Ports** tab.
 - o In the System topic, select **Action > System Settings**, then click the **Ports** tab.
 - o In the Welcome panel, select **System Settings**, and then click the **Ports** tab.
2. From the Host Post Mode dropdown, select **iSCSI**.
3. From the Port Settings tab, set the port-specific options:
 - o 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 3: 10.10.10.100
 - Controller A port 4: 10.11.10.120
 - Controller B port 3: 10.10.10.110
 - Controller B port 4: 10.11.10.130
 - o Netmask. For IPv4, subnet mask for assigned port IP address. The default is 255.255.255.0.
 - o Gateway. For IPv4, gateway IP address for assigned port IP address. The default is 0.0.0.0.
 - o Default Router. For IPv6, default router for 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).
4. From the Advanced Settings tab, set the options that apply to all iSCSI ports:
 - o Enable Authentication (CHAP). Enables or disables use of Challenge Handshake Authentication Protocol. Enabling or disabling CHAP in this panel will update its setting in the Configure CHAP panel (available in the Hosts topic by selecting **Action > Configure CHAP**). Disabled by default.

NOTE: CHAP records for iSCSI login authentication must be defined if CHAP is enabled. To create CHAP records, see [“Configuring CHAP” \(page 84\)](#).

- o Enable Jumbo Frames. Enables or disables support for jumbo frames. Allowing for 100 bytes of overhead, a normal frame can contain a 1400-byte payload whereas a jumbo frame can contain a maximum 8900-byte payload for larger data transfers. Disabled by default.

NOTE: Use of jumbo frames can succeed only if jumbo-frame support is enabled on all network components in the data path.

- o iSCSI IP Version. Specifies whether IP values use Internet Protocol version 4 (IPv4) or version 6 (IPv6) format. IPv4 uses 32-bit addresses. IPv6 uses 128-bit addresses. The default is IPv4.
- o Enable iSNS. Enables or disables registration with a specified Internet Storage Name Service server, which provides name-to-IP-address mapping. Disabled by default.

- o iSNS Address. Specifies the IP address of an iSNS server. The default address is all zeroes.
- o Alternate iSNS Address. Specifies the IP address of an alternate iSNS server, which can be on a different subnet. The default address is all zeroes.

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

5. Perform one of the following:
 - o To save your settings and continue configuring your system, click **Apply**.
 - o To save your settings and close the panel, click **Apply and Close**.

A confirmation panel displays.

6. Click **OK** to save your changes. Otherwise, click **Cancel**.

To configure two ports as FC and two ports as iSCSI per controller

1. Perform one of the following to access the options in the Ports tab:
 - o In the Home topic, select **Action > System Settings**, then click the **Ports** tab.
 - o In the System topic, select **Action > System Settings**, then click the **Ports** tab.
 - o In the Welcome panel, select **System Settings**, and then click the **Ports** tab.
2. From the Host Post Mode dropdown, select **FC-and-iSCSI**.
3. From the Port Settings tab, set the FC port-specific options:
 - o Set the Speed option to the proper value to communicate with the host, or to auto, which auto-negotiates the proper link speed. Because a speed mismatch prevents communication between the port and host, set a speed only if you need to force the port to use a known speed. The maximum link speed is determined by the installed SFP.
 - o Set The FC Connection Mode to either point-to-point or auto:
 - **point-to-point:** Fibre Channel point-to-point. This is the default.
 - **auto:** Automatically sets the mode based on the detected connection type.
4. Set the port-specific options:
 - o 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 3: 10.10.10.100
 - Controller A port 4: 10.11.10.120
 - Controller B port 3: 10.10.10.110
 - Controller B port 4: 10.11.10.130
 - o Netmask. For IPv4, subnet mask for assigned port IP address. The default is 255.255.255.0.
 - o Gateway. For IPv4, gateway IP address for assigned port IP address. The default is 0.0.0.0.
 - o Default Router. For IPv6, default router for 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).
5. In the Advanced Settings section of the panel, set the options that apply to all iSCSI ports:
 - o Enable Authentication (CHAP). Enables or disables use of Challenge Handshake Authentication Protocol. Enabling or disabling CHAP in this panel will update its setting in the Configure CHAP panel (available in the Hosts topic by selecting **Action > Configure CHAP**). Disabled by default.

NOTE: CHAP records for iSCSI login authentication must be defined if CHAP is enabled. To create CHAP records, see “Configuring CHAP” (page 87).

- Enable Jumbo Frames. Enables or disables support for jumbo frames. Allowing for 100 bytes of overhead, a normal frame can contain a 1400-byte payload whereas a jumbo frame can contain a maximum 8900-byte payload for larger data transfers. Disabled by default.

NOTE: Use of jumbo frames can succeed only if jumbo-frame support is enabled on all network components in the data path.

- iSCSI IP Version. Specifies whether IP values use Internet Protocol version 4 (IPv4) or version 6 (IPv6) format. IPv4 uses 32-bit addresses. IPv6 uses 128-bit addresses. The default is IPv4.
- Enable iSNS. Enables or disables registration with a specified Internet Storage Name Service server, which provides name-to-IP-address mapping. Disabled by default.
- iSNS Address. Specifies the IP address of an iSNS server. The default address is all zeroes.
- Alternate iSNS Address. Specifies the IP address of an alternate iSNS server, which can be on a different subnet. The default address is all zeroes.

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

6. Perform one of the following:
 - To save your settings and continue configuring your system, click **Apply**.
 - To save your settings and close the panel, click **Apply and Close**.A confirmation panel displays.
7. Click **OK** to save your changes. Otherwise, click **Cancel**.

Managing scheduled tasks

The Manage Schedules action is enabled when at least one scheduled task exists. When accessed, you can modify or delete scheduled tasks to:

- Create snapshots
- Reset snapshots
- Run replications

To modify a schedule from the Home topic

1. In the Home topic, select **Action > Manage Schedules**.
2. Select the schedule to modify. The schedule's settings appear at the bottom of the panel.
3. Optional: If you want to replicate the last snapshot in the primary volume, select the **Last Snapshot** check box. At the time of the replication, the snapshot must exist. This snapshot may have been created either manually or by a scheduling the snapshot.

NOTE: This option is unavailable when replicating volume groups.

4. Specify a date and a time in the future to be the first instance when the scheduled task will run, and to be the starting point for any specified recurrence.
 - To set the **Date** value, enter the current date in the format **YYYY-MM-DD**.
 - To set the **Time** value, enter two-digit values for the hour and minutes and select either **AM**, **PM**, or **24H** (24-hour clock).

5. If you want the task to run more than once, select the **Repeat** check box.
 - o Specify how often the task should repeat. Enter a number and select the appropriate time unit. Replications can recur no less than 30 minutes apart.
 - o Either make sure the **End** check box is cleared, which allows the schedule to run without an end date, or select the check box and specify when the schedule should stop running.
 - o Either make sure the **Time Constraint** check box is cleared, which allows the schedule to run at any time, or select the check box to specify a time range within which the schedule should run.
 - o Either make sure the **Date Constraint** check box is cleared, which allows the schedule to run on any day, or select the check box to specify the days when the schedule should run.
6. Click **Apply**. A confirmation panel appears.
7. Click **OK** to continue. Otherwise click **Cancel**. If you clicked OK, the schedule is modified.
8. Click **OK**.

To delete a schedule from the Home topic

1. In the Home topic, select **Action > Manage Schedules**. The Manage Schedules panel opens.
2. Select the schedule to delete.
3. Click **Delete Schedule**. A confirmation panel appears.
4. Click **OK** to continue. Otherwise, click **Cancel**. If you clicked OK, the schedule was deleted.
5. Click **OK**.

3 Working in the System topic

Viewing system components

The System topic enables you to see information about each enclosure and its physical components in front, rear, and tabular views. Components vary by enclosure model.

Front view

The Front tab shows the front of all enclosures in a graphical view. For each enclosure, the front view shows the enclosure ID and other information.

To see more information about an enclosure or disks, hover the cursor over an enclosure ear or a disk. To illuminate a locator LED for an enclosure or disk, select one or more component and click **Turn On LEDs**. To turn off locator LEDs, select one or more component and click **Turn Off LEDs**.

Enclosure Information	ID, status, vendor, model, disk count, WWN, midplane serial number, revision, part number, manufacturing date, manufacturing location, EMP A revision, EMP B revision, EMP A bus ID, EMP B bus ID, EMP A target ID, EMP B target ID, midplane type, enclosure power (watts), PCIe 2-capable, health
Disk Information	Location, serial number, usage, type, size, status, RPM (spinning disk only), SSD life left, manufacturer, model, revision, power on hours, FDE state (for MSA 2050 only), FDE lock key (for MSA 2050 only), job running, sector format, transfer rate, SMART, drive spin down count, health

If a component's health is not OK, the health reason, recommended action, and unhealthy subcomponents are shown to help you resolve problems.

NOTE: Following is more information for selected Disk Information panel items:

- *Power On Hours* refers to the total number of hours that the disk has been powered on since it was manufactured. This value is updated in 30-minute increments.
 - For MSA 2050: *FDE State* refers to the FDE state of the disk. For more information about FDE states, see the CLI Reference Guide.
 - For MSA 2050: *FDE lock keys* are generated from the FDE passphrase and manage locking and unlocking the FDE-capable disks in the system. Clearing the lock keys and power cycling the system denies access to data on the disks.
-

Rear view

The Rear tab shows the rear of all enclosures in a graphical view. The rear view shows enclosure IDs and the presence or absence of power supplies, controller modules, and expansion modules. It also shows controller module IDs, host port types and names, network port IP addresses, and expansion port names.

To see more information, hover the cursor over an enclosure ear or a component. To illuminate a locator LED for any of the components, select one or more components and click **Turn On LEDs**. To turn off locator LEDs, select one or more components and click **Turn Off LEDs**.

NOTE: Protocol-specific properties are displayed only for host ports that use those protocols.

Table 10 Additional information for rear view of enclosure

Enclosure	ID, status, vendor, model, disk count, WWN, midplane serial number, revision, part number, manufacturing date, manufacturing location, EMP A revision, EMP B revision, EMP A bus ID, EMP B bus ID, EMP A target ID, EMP B target ID, midplane type, enclosure power (watts), PCIe 2-capable, health
Power supply	Status, vendor, model, serial number, revision, location, part number, manufacturing date, manufacturing location, health
Controller module	ID, IP address, description, status, model, serial number, hardware version, system cache memory (MB), revision, CPLD version, Storage Controller code version, Storage Controller CPU type, part number, position, hardware version, manufacturing date, manufacturing location, health
FC host port	Name, type, ID (WWN), status, configured speed, actual speed, topology, primary loop ID, supported speeds, SFP status, part number, health
iSCSI host port	Name, type, ID (IQN), status, actual speed, IP version, address, gateway, netmask, supported speeds, SFP status, part number, configured speed, 10G compliance, cable length, cable technology, Ethernet compliance, health
SAS host port	Name, type, ID (WWN), status, configured speed, actual speed, cable type, health
Network port	Name, mode, IP address, network mask, gateway, MAC address, health
Expansion port	Enclosure ID, controller ID, name, status, health
Expansion module (IOM)	ID, description, serial number, hardware revision, health

If a component's health is not OK, the health reason, recommended action, and unhealthy subcomponents are shown to help you resolve problems.

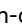





For MSA 1050: If the system is configured to use fan-out cables, fan-out cable icons  appear between the depicted SAS ports. The number of SAS ports that display depends on the configuration.

Table view

The Table tab shows a tabular view of information about physical components in the system. By default, the table shows 20 entries at a time. For information about using tables, see [“Tips for using tables” \(page 12\)](#).

For each component, the table shows the following information:

- Health. Shows the health of the component:  OK,  Degraded,  Fault,  N/A, or  Unknown.
- Type. Shows the component type: enclosure, disk, power supply, controller module, network port, host port, expansion port, CompactFlash card, expander, or I/O module (expansion module).
- Enclosure. Shows the enclosure ID.
- Location. Shows the location of the component.
 - For an enclosure, the location is shown in the format *Rack rack-ID.shelf-ID*. You can set the location through the CLI `set enclosure` command.
 - For a disk, the location is shown in the format *enclosure-ID.disk-slot*.
 - For a power supply or I/O module, the locations Left and Right are as viewed from the rear of the enclosure.
 - For a host port, the location is shown as controller ID and port number.
- Information. Shows additional, component-specific information:
 - For an enclosure: its FRU description and current disk count.
 - For a disk: its type, capacity, and usage.

- Type is shown as either:
 - MDL. Spinning midline SAS disk.
 - SAS. Spinning enterprise-class SAS disk.
 - SSD. Solid-state disk.
- Usage is shown as either:
 - AVAIL. The disk is available.
 - GLOBAL SP. The disk is configured as a global spare.
 - pool-ID:tier name for disk groups that are part of a virtual pool. The disk is part of a disk group.
 - FAILED. The disk is unusable and must be replaced. Reasons for this status include: excessive media errors, SMART error, disk hardware failure, or unsupported disk.
 - LEFTOVR. The disk is part of a disk group that is not found in the system.
 - UNUSABLE. The disk cannot be used in a disk group. Possible reasons include: the system is secured and the disk is data locked; the system is secured/locked (no passphrase available) and the disk is data locked; the system is secured and the disk is not FDE.
- For a power supply: its FRU description.
- For a controller module: its ID.
- For a network port: its IP address.
- For a host port: one of the following values:
 - FC(L). Fibre Channel-Arbitrated Loop (public or private)
 - FC(P). Fibre Channel Point-to-Point
 - FC(-). Fibre Channel disconnected
 - SAS. Serial Attached SCSI
 - iSCSI. Internet SCSI
- For an expansion port: either Out Port or In Port.
- For an I/O module: its ID.
- Status. Shows the component status:
 - For an enclosure: Up.
 - For a disk:
 - Up. The disk is present and is properly communicating with the expander.
 - Spun Down. The disk is present and has been spun down by the DSD feature.
 - Warning. The disk is present but the system is having communication problems with the disk LED processor. For disk and midplane types where this processor also controls power to the disk, power-on failure will result in the Error status.
 - Error. The disk is present but not detected by the expander.
 - Unknown. Initial status when the disk is first detected or powered on.
 - Not Present. The disk slot indicates that no disk is present.
 - Unrecoverable. The disk is present but has unrecoverable errors.
 - Unavailable. The disk is present but cannot communicate with the expander.
 - Unsupported. The disk is present but is an unsupported type.
 - For a power supply: Up, Warning, Error, Not Present, or Unknown.
 - For a controller module or I/O module: Operational, Down, Not Installed, or Unknown.
 - For a network port: N/A.

- For a host port:
 - Up. The port is cabled and has an I/O link.
 - Warning. Not all of the port's PHYs are up.
 - Error. The port is reporting an error condition.
 - Not Present. The controller module is not installed or is down.
 - Disconnected. Either no I/O link is detected or the port is not cabled.
- For an expansion port: Up, Disconnected, or Unknown.
- For a CompactFlash card: Installed, Not Installed, or Unknown.

Configuring system settings

The System Settings panel provides options for you to quickly and easily configure your system. Access the panel by doing one of the following:

- In the Home topic, select **Action > System Settings**.
- In the System topic, select **Action > System Settings**.
- In the Welcome panel, select **System Settings**.

For more information on configuring system setting options, see [“Configuring system settings” \(page 64\)](#).

Resetting host ports

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 fix such a problem you might need to reset controller host ports (channels).

For FC, you can reset a single port. For an FC host port configured to use FC-AL (loop) topology, a reset issues a loop initialization primitive (LIP).

For iSCSI, you can reset a port pair (either the first and second ports or the third and fourth ports).

For SAS, you can reset a port pair. Resetting a SAS host port issues a COMINT/COMRESET sequence and might reset other ports.

To reset a host port

1. In the **System** topic, select **Action > Reset Host Port**.
2. Select the port or port pair to reset.
3. Click **OK** and follow the prompts.

Rescanning disk channels

A rescan forces a rediscovery of disks and enclosures in the storage system. If both Storage Controllers are online and can communicate with both expansion modules in each connected enclosure, a rescan also reassigns enclosure IDs to follow the enclosure cabling order of controller A. For further cabling information, refer to your product's User 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.

To rescan disk channels

1. Verify that both controllers are operating normally.

2. Perform one of the following:
 - o Point to the **System** tab and select **Rescan Disk Channels**.
 - o In the System topic, select **Action > Rescan Disk Channels**.
The Rescan Disk Channels panel opens.
3. Click **Rescan**.

Clearing disk metadata

You can clear metadata from a leftover disk to make it available for use.

CAUTION:

- o Only use this command when all disk groups are online and leftover disks exist. Improper use of this command may result in data loss.
- o Do not use this command when a disk group is offline and one or more leftover disks exist.
- o If you are uncertain whether to use this command, 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 a *leftover*:

- 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 that is a member of a disk group in another system is moved into this system without the other members of its group.

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 spare is available, and the health of the disk group is Degraded or Critical, the disk group will use them to start reconstruction. When reconstruction is complete, you can clear the leftover disk's metadata. Clearing the metadata will 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.

This command clears metadata from leftover disks only. If you specify disks that are not leftovers, the disks are not changed.

To clear metadata from leftover disks

1. In the System topic, select **Action > Clear Metadata**. The Clear Metadata panel opens.
2. Select the leftover disks from which to clear metadata.
3. Click **OK**.
4. Click **Yes** to continue. Otherwise, click **No**. If you clicked Yes, the metadata is cleared.
5. Click **OK**.

Updating firmware

You can view the current versions of firmware in controller modules, expansion modules, and disk drives. As a user with the `manage` role, you can also install new versions. For information about supported releases for firmware update, see the Release Notes for your product. For information about which controller module will update the other when a controller module is replaced, see [“About firmware update” \(page 30\)](#).

To monitor the progress of a firmware-update operation by using the activity progress interface, see [“Using the activity progress interface” \(page 69\)](#).

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

Best practices for firmware update

- In the health panel in the footer, verify that the system health is OK. If the system health is not OK, view the Health Reason value in the health panel in the footer and resolve all problems before you update firmware. For information about the health panel, see [“Viewing health information” \(page 141\)](#).
- 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 need to 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.
- If a disk group is quarantined, resolve the problem that is causing the component 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 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.

Updating controller module firmware

In a dual-controller system, both controller modules should run the same firmware version. Storage systems in a replication set should run the same or compatible firmware versions. You can update disk-drive firmware by searching on <http://www.hpe.com/storage/msadrivefirmware> for your drive model number to find the latest firmware to download. Then, load the firmware file obtained from the HPE web download site at <http://www.hpe.com/support/hpesc>. To install an HPE ROM Flash Component or firmware Smart Component, follow the instructions on the HPE web site. Otherwise, to install a firmware binary file, follow the steps below.

To prepare to update controller module firmware

1. Follow the best practices in [“Best practices for firmware update” \(page 66\)](#).
2. Obtain the appropriate firmware file and download it to your computer or network.

To update controller module firmware

1. As a user with the `manage` role, perform one of the following:
 - In the System topic, select **Action > Update Firmware**.
 - In the banner, click the system panel and select **Update Firmware**.
 - In the Welcome panel, select **Upgrade Firmware**.The Update Firmware panel opens. The Update Controller Modules tab shows versions of firmware components that are currently installed in each controller.
2. Click the **Bundle or Controller Firmware File** button to select the firmware file to install.

3. Optional: Select (enable) or clear (disable) the PFU check box and confirm the action.

NOTE: For information about which controller module will update the other when a controller module is replaced, see “About firmware update” (page 30).

4. Click **OK**. A panel shows firmware-update progress.

The process starts by validating the firmware file:

- If the file is invalid, verify that you specified the correct firmware file. If you did, try downloading it again from the source location.
- If the file is valid, the process continues.

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.

Firmware update typically takes 10 minutes for a controller with current CPLD firmware, or 20 minutes for a controller with downlevel CPLD 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 each EMP in an MSA 1050/2050 drive 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. If this problem persists, contact technical support.

When firmware update on the local controller is complete, users are automatically signed out and the MC 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, allow an additional 10–20 minutes for the partner controller to be updated.

5. Clear your web browser cache, then sign in to the SMU. If PFU is running on the controller you sign in to, a panel shows PFU progress and prevents you from performing other tasks until PFU is complete.

NOTE: If PFU is enabled for the system, after firmware update has completed on both controllers, check the system health. If the system health is Degraded and the health reason indicates that the firmware version is incorrect, verify that you specified the correct firmware file 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 enclosure management processor (EMP). All modules of the same model should run the same firmware version.

You can update disk-drive firmware by searching on <http://www.hpe.com/storage/msadrivefirmware> for your drive model number to find the latest firmware to download. Then, load the firmware file obtained from the HPE web download site at <http://www.hpe.com/support/hpesc>. To install an HPE ROM Flash Component or firmware Smart Component, follow the instructions on the HPE web site. Otherwise, to install a firmware binary file, follow the steps below.

To prepare to update expansion module firmware

1. Follow the best practices in “Best practices for firmware update” (page 66).
2. Obtain the appropriate firmware file and download it to your computer or network.

To update expansion module firmware

1. As a user with the `manage` role, perform one of the following:
 - o In the System topic, select **Action > Update Firmware**.
 - o In the banner, click the system panel and select **Update Firmware**.
 - o In the Welcome panel, select **Upgrade Firmware**.
The Update Firmware panel opens.
2. Select the **Update Expansion Modules** tab. This tab shows information about each expansion module in the system.
3. Select the expansion modules to update.
4. Click **File** and select the firmware file to install.
5. Click **OK**. Messages show firmware update progress.

△ CAUTION: Do not perform a power cycle or controller restart during the firmware update. If the update is interrupted or there is a power failure, the module might become inoperative. If this occurs, contact technical support. The module might need to be returned to the factory for reprogramming.

It typically takes 2.5 minutes to update each EMP in an MSA 1050/2050 drive enclosure.

6. Verify that each updated expansion module has the new firmware version.

Updating disk-drive firmware

You can update disk-drive firmware by searching on <http://www.hpe.com/storage/msadrivefirmware> for your drive model number to find the latest firmware to download. Then, load the firmware file obtained from the HPE web download site at <http://www.hpe.com/support/hpesc>. To install an HPE ROM Flash Component or firmware Smart Component, follow the instructions on the HPE web site. Otherwise, to install a firmware binary file, follow the steps below.

A dual-ported disk drive can be updated from either controller.

To prepare to update disk-drive firmware

1. Follow the best practices in “[Best practices for firmware update](#)” (page 66).
2. Obtain the appropriate firmware file and download it to your computer or network.
3. Stop I/O to the storage system. During the update all volumes will be temporarily inaccessible to hosts. If I/O is not stopped, mapped hosts will report I/O errors. Volume access is restored after the update completes.

To update disk-drive firmware

1. As a user with the `manage` role, perform one of the following:
 - o In the System topic, select **Action > Update Firmware**.
 - o In the banner, click the system panel and select **Update Firmware**.
 - o In the Welcome panel, select **Upgrade Firmware**.
The Update Firmware panel opens.
2. Select the **Update Disk Drives** tab. This tab shows information about each disk drive in the system.
3. Select the disk drives to update.
4. Click **File** and select the firmware file to install.

5. Click OK.

⚠ CAUTION: Do not power cycle enclosures or restart a controller during the firmware update. If the update is interrupted or there is a power failure, the disk drive might become inoperative. If this occurs, contact technical support.

It typically takes several minutes for the firmware to load. Wait for a message that the update has completed.

6. Verify that each disk drive has the new firmware revision.

Using the activity progress interface

The activity progress interface reports whether a firmware update or partner firmware update operation is active and shows the progress through each step of the operation. In addition, when the update operation completes, status is presented indicating either the successful completion, or an error indication if the operation failed.

To use the activity progress interface

1. Enable the Activity Progress Monitor service. See [“Enabling or disabling system-management services” \(page 50\)](#).
2. In a new tab in your web browser, enter the URL for the form:

`http://controller-address:8081/cgi-bin/content.cgi?mc=MC-identifier&refresh=true`
where:

- *controller-address* is required and specifies the IP address of a controller network port.
- *mc=MC-identifier* is an optional parameter that specifies the controller for which to report progress/status:
 - *mc=A* shows output for controller A only.
 - *mc=B* shows output for controller B only.
 - *mc=both* shows output for both controllers.
 - *mc=self* shows output for the controller whose IP address is specified.
- *refresh=true* is an optional parameter that causes automatic refresh of the displayed output every second. This will continue until either:
 - The parameter is removed.
 - The controller whose IP address is specified is restarted and communication is lost.

When activity is in progress, the interface will display an MC-specific Activity Progress table with the following properties and values.

Table 11 Activity progress properties and values

Property	Value
Time	The date and time of the latest status update.
Seconds	The number of seconds this component has been active.
Component	The name of the object being processed.
Status	The status of the component representing its progress/completion state. <ul style="list-style-type: none"> • ACTIVE: The operation for this component is currently active and in progress. • OK: The operation for this component completed successfully and is now inactive. • N/A: The operation for this component was not completed because it was not applicable. • ERROR: The operation for this component failed with an error (see code and message).

Table 11 Activity progress properties and values

Property	Value
Code	A numeric code indicating the status. <ul style="list-style-type: none">• 0: The operation for this component completed with a “completed successfully” status.• 1: The operation for this component was not attempted because it is not applicable (the component doesn’t exist or doesn’t need updating).• 2: The operation is in progress. The other properties will indicate the progress item (message, current, total, percent).• 10 or higher: The operation for this component completed with a failure. The code and message indicate the reason for the error.
Message	A textual message indicating the progress status or error condition.

Changing FDE settings (for MSA 2050 only)

In the Full Disk Encryption panel, you can change settings for these options:

- FDE general configuration
 - Set the passphrase
 - Clear lock keys
 - Secure the system
 - Repurpose the system
- Repurpose disks
- Set import lock key IDs

Changing FDE general configuration

⚠ CAUTION: Do not change FDE configuration settings while running I/O. Temporary data unavailability may result. Also, the intended configuration change might not take effect.

Setting the passphrase

You can set the FDE passphrase the system uses to write to and read from FDE-capable disks. From the passphrase, the system generates the lock key ID that is used to secure the FDE-capable disks. If the passphrase for a system is different from the passphrase associated with a disk, the system cannot access data on the disks.

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

To set or change the passphrase

1. In the System topic, select **Action > Full Disk Encryption**.
The Full Disk Encryption panel opens with the **FDE General Configuration** tab selected.
2. Enter a passphrase in the **Passphrase** field of the **Set/Create Passphrase** section. A passphrase is case sensitive and can include 8–32 printable UTF-8 characters except for the following: , < > \
3. Re-enter the passphrase.

4. Perform one of the following:

- To secure the system now, click the **Secure** checkbox, then select **Set and Secure**. A dialog box will confirm the passphrase was changed successfully.
- To save the passphrase without securing the system, click **Set**. A dialog box will confirm the passphrase was changed successfully. To secure the system at a later date, see [“Securing the system” \(page 71\)](#).

Clearing lock keys

Lock keys are generated from the passphrase and manage locking and unlocking the FDE-capable disks in the system. Clearing the lock keys and power cycling the system denies access to data on the disks. Use this procedure 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; disk group status will become QTOF; pool health will become Degraded; and volumes will become inaccessible.

To restore access to 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.

To clear lock keys

NOTE: The FDE tabs are dynamic, and the **Clear All FDE Keys** option is not available on a secured system until the current passphrase is entered in the Current Passphrase field. (If you do not have a passphrase, the **Clear All FDE Keys** option will not appear. If you have a passphrase but have not entered it, you can view, but will be unable to access, this option.) If there is no passphrase, set one using the procedure in [“Setting the passphrase” \(page 70\)](#).

1. In the System topic, select **Action > Full Disk Encryption**.

The Full Disk Encryption panel opens with the **FDE General Configuration** tab selected.

2. Enter the passphrase in the Current Passphrase field.

3. Click **Clear**. A dialog box displays.

4. At the confirmation prompt, perform one of the following:

- To clear the lock keys for the system, click **OK**.
- To cancel the request, click **Cancel**.

Securing the system

An FDE-capable system must be secured to enable FDE protection.

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

To secure the system

NOTE: The FDE tabs are dynamic, and the **Secure** option is not available until the current passphrase is entered in the Current Passphrase field. (If you do not have a passphrase, the **Secure** option will not appear. If you have a passphrase but have not entered it, you can view but will be unable to access this option.) If there is no passphrase, set one using the procedure in [“Setting the passphrase” \(page 70\)](#).

1. In the System topic, select **Action > Full Disk Encryption**.

The Full Disk Encryption panel opens with the **FDE General Configuration** tab selected.

2. Enter the passphrase in the Current Passphrase field.

3. Click **Secure**. A message displays confirming that the system is in a secure state.

Repurposing the system

You can repurpose a system to erase all data on the system and return its FDE state to unsecure.

CAUTION: Repurposing a system erases all disks in the system and restores the FDE state to unsecure.

To repurpose the system

NOTE: The FDE tabs are dynamic, and the **Repurpose System** option is not available until the system is secure and all disk groups have been removed from the system.

1. Delete all disk groups in the system. To delete disk groups, see [“Removing disk groups” \(page 91\)](#). Removing disk groups effectively deletes all data on the disks but does not secure erase them.
2. Click the **System** tab.
3. In the System topic, select **Action > Full Disk Encryption**.
The Full Disk Encryption panel opens with the **FDE General Configuration** tab selected.
4. In the Repurpose System section, click the **Repurpose** button.
5. At the confirmation prompt, perform one of the following:
 - o To repurpose the system, click **OK**.
 - o To cancel the request, click **Cancel**.

Repurposing disks

You can 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 unsecure 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.

To repurpose a disk

1. In the System topic, select **Action > Full Disk Encryption**.
The Full Disk Encryption panel opens with the **FDE General Configuration** tab selected.
2. Select the **Repurpose Disks** tab.
3. Perform one of the following:
 - o Select the disks to repurpose, then choose **Repurpose** and follow the confirmation prompts.
 - o Check **Select all** to repurpose all FDE disks in the system, then click **Repurpose** and follow the confirmation prompts.
 - o To cancel the request, click **Cancel**.

Setting import lock key IDs

You can set the passphrase associated with an import lock key to unlock FDE-secured disks that are inserted into the system from a different secure system. If the correct passphrase is not entered, the system cannot access data on the disk.

After importing disks into the system, the disks will now be associated with the system lock key ID and data will no longer be accessible using the import lock key. This effectively transfers security to the local system passphrase.

To set or change the import passphrase

1. In the System topic, select **Action > Full Disk Encryption**.
The Full Disk Encryption panel opens with the **FDE General Configuration** tab selected.
2. Select the **Set Import Lock Key ID** tab.
3. In the Passphrase field, enter the passphrase associated with the displayed lock key.
4. Re-enter the passphrase.
5. Click **Set**. A dialog box will confirm the passphrase was changed successfully.

Configuring advanced settings

Use the Advanced Settings panel to change disk settings, cache settings, partner firmware update settings, and system utility settings.

Changing disk settings

The Disk tab provides options to change disk settings, including SMART configuration, EMP polling rate, dynamic spares, and drive spin down options.

Configuring SMART

Self-Monitoring Analysis and Reporting Technology (SMART) provides data that enables you to monitor disks and analyze why a disk failed. When SMART is enabled, the system checks for SMART events one minute after a restart and every five minutes thereafter. SMART events are recorded in the event log.

To change the SMART setting

1. In the System topic, select **Action > Advanced Settings > Disk**.
2. Set the SMART Configuration option to one of the following:
 - o **Don't Modify**. Allows current disks to retain their individual SMART settings and does not change the setting for new disks added to the system.
 - o **Enabled**. Enables SMART for all current disks after the next rescan and automatically enables SMART for new disks added to the system. This option is the default.
 - o **Disabled**. Disables SMART for all current disks after the next rescan and automatically disables SMART for new disks added to the system.
3. Click **Apply**. If you chose to disable SMART, a confirmation panel displays. Click **Apply** to accept the changes or click **Cancel**.

Configuring the EMP polling rate

You can change the frequency interval that the storage system polls each attached enclosure's management processor (EMP) for changes to temperature, power supply and fan status, and the presence or absence of disks. Typically you can use the default setting.

- Increasing the interval might slightly improve processing efficiency, but changes in device status are communicated less frequently. For example, this increases the amount of time before LEDs are updated to reflect status changes.
- Decreasing the interval slightly decreases processing efficiency, but changes in device status are communicated more frequently. For example, this decreases the amount of time before LEDs are updated to reflect status changes.

To change the EMP polling rate

1. In the System topic, select **Action > Advanced Settings > Disk**.
2. Set the EMP Polling Rate interval. The options are 5, 10, or 30 seconds; or 1, 5, 10, 15, 20, 25, 30, 45, or 60 minutes. The default is 5 seconds.
3. Click **Apply**.

Configuring dynamic spares

The dynamic spares feature lets you use all of your disks in fault-tolerant disk groups without designating a disk as a spare. With dynamic spares enabled, if a disk fails and you replace it with a compatible disk, the storage system rescans the bus, finds the new disk, automatically designates it a spare, and starts reconstructing the disk group. A compatible disk has enough capacity to replace the failed disk and is the same type (as described in [“About spares” \(page 22\)](#)). If a spare or available compatible disk is already present, the dynamic spares feature uses that disk to start the reconstruction and the replacement disk can be used for another purpose.

To change the dynamic spares setting

1. In the System topic, select **Action > Advanced Settings > Disk**.
2. Either select (enable) or clear (disable) the **Dynamic Spare Capability** option. Enabled by default.
3. Click **Apply**. If you chose to disable dynamic spares, a confirmation panel displays. Click **Apply** to accept the changes or click **Cancel**.

Configuring drive spin down for available disks and global spares

For spinning disks, the drive spin down (DSD) feature monitors disk activity within system enclosures and spins down inactive disks to conserve energy. You can enable or disable DSD for available spinning disks that are not in a virtual pool, and for global spares. You can also set the period of inactivity after which available disks and global spares automatically spin down.

To configure a time period to suspend and resume DSD for all disks, see [“Scheduling drive spin down for available disks and global spares” \(page 74\)](#).

DSD affects disk operations as follows:

- Spun-down disks are not polled for SMART events.
- Operations requiring access to disks may be delayed while the disks are spinning back up.

To configure DSD for available disks and global spares

1. In the System topic, select **Action > Advanced Settings > Disk**.
2. Set the options:
 - o Either select (enable) or clear (disable) the **Available and Spare Drive Spin Down Capability** option.
 - o Set the **Drive Spin Down Delay (minutes)** option, which is the period of inactivity after which available disks and global spares automatically spin down, from 1–360 minutes. The default is 15 minutes.
3. Click **Apply**. When processing is complete a success dialog appears.
4. Click **OK**.

Scheduling drive spin down for available disks and global spares

For all spinning disks that are configured to use drive spin down (DSD), you can configure a time period to suspend and resume DSD so that disks remain spun-up during hours of frequent activity.

To configure DSD for available disks and global spares, see [“Configuring drive spin down for available disks and global spares” \(page 74\)](#).

DSD affects disk operations as follows:

- Spun-down disks are not polled for SMART events.
- Operations requiring access to disks may be delayed while the disks are spinning back up.
- If a suspend period is configured and it starts while a disk has started spinning down, the disk spins up again.

To schedule DSD for all spinning disks

1. In the System topic, select **Action > Advanced Settings > Disk**.
2. Set the options:
 - o Select the **Drive Spin Down Suspend Period** option.
 - o Set the **Time to Suspend** and **Time to Resume** options. For each, enter hour and minutes values and select either **AM**, **PM**, or **24H** (24-hour clock).
 - o If you want the schedule to apply only Monday through Friday, select the **Exclude Weekend Days from Suspend Period** option.
3. Click **Apply**. When processing is complete a success dialog appears.
4. Click **OK**.

Changing system cache settings

The Cache tab provides options to change the synchronize-cache mode, missing LUN response, host control of the system's write-back cache setting, and auto-write-through cache triggers and behaviors.

Changing the synchronize-cache mode

You can control how the storage system handles the SCSI `SYNCHRONIZE CACHE` command. Typically you can use the default setting. However, if the system has performance problems or problems writing to databases or other applications, contact technical support to determine if you should change this option.

To change the synchronize-cache mode

1. In the System topic, select **Action > Advanced Settings > Cache**.
2. Set the Sync Cache Mode option to either:
 - o **Immediate**. Good status is returned immediately and cache content is unchanged. This is the default.
 - o **Flush to Disk**. Good status is returned only after all write-back data for the specified volume is flushed to disk.
3. Click **Apply**.

Changing the missing LUN response

Some operating systems do not look beyond LUN 0 if they do not find a LUN 0 or cannot handle noncontiguous LUNs. The Missing LUN Response option handles these situations by enabling the host drivers to continue probing for LUNs until they reach the LUN to which they have access.

This option controls the SCSI sense data returned for volumes that are not accessible because they don't exist or have been hidden through volume mapping (this does not apply to volumes of offline disk groups). Use the default value, Not Ready, unless the system is used in a VMware environment or a service technician asks you to change it to work around a host driver problem.

To change the missing LUN response

1. In the System topic, select **Action > Advanced Settings > Cache**.
2. Set the Missing LUN Response option to either:
 - o **Not Ready**. Sends a reply that there is a LUN where a gap has been created but that it's "not ready." Sense data returned is a Sense Key of 2h and an ASC/ASCQ of 04/03. This option is the default.
 - o **Illegal Request**. Sends a reply that there is a LUN but that the request is "illegal." Sense data returned is a Sense Key of 5h and an ASC/ASCQ of 25/00. If the system is used in a VMware environment, use this option.
3. Click **Apply**.

Controlling host access to the system's write-back cache setting

You can prevent hosts from using SCSI `MODE SELECT` commands to change the system's write-back cache setting. Some operating systems disable write cache. If host control of write-back cache is disabled, the host cannot modify the cache setting. The default is Disabled.

This option is useful in some environments where the host disables the system's write-back cache, resulting in degraded performance.

To change host access to the write-back cache setting

1. In the System topic, select **Action > Advanced Settings > Cache**.
2. Either select (enable) or clear (disable) the **Host Control of Write-Back Cache** option.
3. Click **Apply**.

Changing auto-write-through cache triggers and behaviors

You can set conditions that cause ("trigger") a controller to change the cache mode from write-back to write-through, as described in ["About volume cache options" \(page 23\)](#). You can also specify actions for the system to take when write-through caching is triggered.

To change auto-write-through cache triggers and behaviors

1. In the System topic, select **Action > Advanced Settings > Cache**.
2. In the Auto-Write Through Cache Trigger Conditions section, either select (enable) or clear (disable) the options:
 - o **Controller Failure**. Changes to write-through if a controller fails. In a dual-controller system this option is disabled by default.
 - o **Cache Power**. Changes to write-through if cache backup power is not fully charged or fails. Enabled by default.
 - o **CompactFlash**. Changes to write-through if CompactFlash memory is not detected during POST, fails during POST, or fails while the controller is under operation. Enabled by default.
 - o **Power Supply Failure**. Changes to write-through if a power supply unit fails. Disabled by default.
 - o **Fan Failure**. Changes to write-through if a cooling fan fails. Disabled by default.
 - o **Overtemperature Failure**. Forces a controller shutdown if a temperature is detected that exceeds system threshold limits. Disabled by default.
3. In the Auto-Write Through Cache Behaviors section, either select (enable) or clear (disable) the options:
 - o **Revert when Trigger Condition Clears**. Changes back to write-back caching after the trigger condition is cleared. Enabled by default.
 - o **Notify Other Controller**. Notifies the partner controller that a trigger condition occurred. Enable this option to have the partner also change to write-through mode for better data protection. Disable this option to allow the partner to continue using its current caching mode for better performance. In a dual-controller system this option is disabled by default.
4. Click **Apply**. If you disabled Cache Power or CompactFlash, a confirmation prompt displays. Choose **Apply** to accept the changes, or **Cancel** to discard the changes.

Configuring partner firmware update

In a dual-controller system in which partner firmware update is enabled (the default), when you update firmware on one controller, the system automatically updates the partner controller. Disable partner firmware update only if requested by a service technician.

To change the partner firmware update setting

1. In the System topic, select **Action > Advanced Settings > Firmware**.
2. Either select (enable) or clear (disable) the **Partner Firmware Update** option.
3. Click **Apply**.

Configuring system utilities


The System Utilities tab lets you configure background scrub for disk groups and individual disks, set utility priority, and enable or disable managed logs.

Configuring background scrub for disk groups

You can enable or disable whether the system continuously analyzes disks in disk groups to find and fix disk errors. This command will fix parity mismatches for RAID 5 and 6; find but not fix mirror mismatches for RAID 1 and 10. It will not fix media errors.

You can use a disk group while it is being scrubbed. Background disk group scrub runs at background utility priority, which reduces to no activity if processor usage is above a certain percentage or if I/O is occurring on the disk group being scrubbed. A disk group scrub may be in process on multiple disk groups at once. A new disk group will first be scrubbed 20 minutes after creation. After a disk group is scrubbed, scrub will start again after the interval specified by the Disk Group Scrub Interval (hours) option.

When a scrub is complete, event 207 is logged and specifies whether errors were found and whether user action is required. Enabling background disk group scrub is recommended.

 **TIP:** If you choose to disable background disk group scrub, you can still scrub a selected disk group. See [“Verifying and scrubbing disk groups” \(page 94\)](#).

To configure background scrub for disk groups

1. In the System topic, choose **Action > Advanced Settings > System Utilities**.
2. Set the options:
 - o Either select (enable) or clear (disable) the **Disk Group Scrub** option. This option is enabled by default.
 - o Set the **Disk Group Scrub Interval (hours)** option, which is the interval between background disk group scrub finishing and starting again, from 0–360 hours. The default is 24 hours.
3. Click **Apply**. If you chose to disable background scrub, a confirmation panel appears. Click **Apply** to accept the changes or click **Cancel**.

Configuring background scrub for disks not in disk groups

You can enable or disable whether the system continuously analyzes disks that are not in disk groups to find and fix disk errors. The interval between background disk scrub finishing and starting again is 72 hours. The first time you enable this option, background disk scrub will start with minimal delay. If you disable and then re-enable this option, background disk scrub will start 72 hours after the last background disk scrub completed.

Enabling background disk scrub is recommended for SAS disks.

To configure background scrub for disks not in disk groups

1. In the System topic, choose **Action > Advanced Settings > System Utilities**.
2. Either select (enable) or clear (disable) the **Disk Scrub** option. This option is disabled by default.
3. Click **Apply**.

Configuring utility priority

You can change the priority at which the Verify, Reconstruct, Expand, and Initialize utilities run when there are active I/O operations competing for the system's controllers.

To change the utility priority

1. In the System panel, choose **Action > Advanced Settings > System Utilities**.

2. Set the Utility Priority option to either:
 - o **High.** Use when your highest priority is to get the system back to a fully fault-tolerant state. This causes heavy I/O with the host to be slower than normal. This value is the default.
 - o **Medium.** Use when you want to balance data streaming with data redundancy.
 - o **Low.** Use when streaming data without interruption, such as for a web server, is more important than data redundancy. This enables a utility such as Reconstruct to run at a slower rate with minimal effect on host I/O.
3. Click **Apply**.

Enabling/disabling managed logs

You can enable or disable the managed logs feature, which allows log files to be transferred from the storage system to a log-collection system to avoid losing diagnostic data. For an overview of the managed logs feature, including how to configure and test it, see [“About managed logs” \(page 30\)](#).

To enable or disable managed logs

1. In the System topic, select **Action > Advanced Settings > System Utilities**.
2. Either select (enable) or clear (disable) the **Managed Logs** option. This option is disabled by default.
3. Click **Apply**.

Restarting or shutting down controllers

Each controller module contains a Management Controller processor and a Storage Controller processor. When necessary, you can restart or shut down these processors for one controller or both controllers.

Restarting controllers

Perform a restart when the SMU informs you that you have changed a configuration setting that requires a restart or when the controller is not working properly.

When you restart a Management Controller, communication with it is lost until it successfully restarts. If the restart fails, the Management Controller in the partner controller module remains active with full ownership of operations and configuration information.

When you restart a Storage Controller, it attempts to shut down with a proper failover sequence. This sequence includes stopping all I/O operations and flushing the write cache to disk. At the end, the controller restarts. Restarting a Storage Controller restarts the corresponding Management Controller.

CAUTION: If you restart both controller modules, all users will lose access to the system and its data until the restart is complete.

NOTE: When a Storage Controller is restarted, current performance statistics that it recorded are reset to zero, but historical performance statistics are not affected. In a dual-controller system, disk statistics may be reduced but are not reset to zero, because disk statistics are shared between the two controllers. For more information, see [“Viewing performance statistics” \(page 135\)](#).

To perform a restart

1. Perform one of the following:
 - o In the banner, click the system panel and select **Restart System**.
 - o In the System topic, select **Action > Restart System**.
The Controller Restart and Shut Down panel opens.
2. Select the **Restart** operation.

3. Select the controller type to restart: **Management** or **Storage**.
4. Select the controller module to restart: **Controller A**, **Controller B**, or both.
5. Click **OK**. A confirmation panel appears.
6. Click **OK** to continue. Otherwise, click **Cancel**. If you clicked **OK**, a message describes restart activity.

Shutting down controllers

Perform a shut down before you remove a controller module from an enclosure, or before you power off its enclosure for maintenance, repair, or a move. Shutting down the Storage Controller in a controller module ensures that a proper failover sequence is used, which includes stopping all I/O operations and writing any data in write cache to disk. If you shut down the Storage Controller in both controller modules, hosts cannot access system data and the SMU will be unavailable.

△ CAUTION: You can continue to use the CLI when either or both Storage Controllers are shut down, but some information might not be available.

To perform a shut down

1. Perform one of the following:
 - In the banner, click the system panel and select **Restart System**.
 - In the System topic, select **Action > Restart System**.
The Controller Restart and Shut Down panel opens.
2. Select the **Shut Down** operation, which automatically selects the Storage controller type.
3. Select the controller module to shut down: **Controller A**, **Controller B**, or both.
4. Click **OK**. A confirmation panel appears.
5. Click **OK** to continue. Otherwise, click **Cancel**. If you clicked **OK**, a message describes shutdown activity.

4 Working in the Hosts topic

Viewing hosts

The Hosts topic shows a tabular view of information about initiators, hosts, and host groups that are defined in the system. For information about using tables, see [“Tips for using tables” \(page 12\)](#). For more information about hosts, see [“About initiators, hosts, and host groups” \(page 25\)](#). The Hosts topic also enables users to map initiators (see [page 111](#)) and view map details (see [page 114](#)).

Hosts table

The hosts table shows the following information. By default, the table shows 10 entries at a time.

- Group. Shows the group name if the initiator is grouped into a host group; otherwise, --.
- Host. Shows the host name if the initiator is grouped into a host; otherwise, --.
- Nickname. Shows the nickname assigned to the initiator.
- ID. Shows the initiator ID, which is the WWN of an FC or SAS initiator or the IQN of an iSCSI initiator.
- Profile. Shows profile settings:
 - Standard. Default profile.
 - HP-UX. The host uses Flat Space Addressing.
 - OpenVMS. LUN 0 cannot be assigned to a mapping.
- Discovered. Shows Yes for a discovered initiator, or No for an initiator that is currently not logged into the system.
- Mapped. Shows Yes for an initiator that is mapped to volumes, or No for an initiator that is not mapped.
- Host Type. Shows the host interface protocol.

Related Maps table

For selected initiators, the Related Maps table shows the following information. By default, the table shows 20 entries at a time.

- Group.Host.Nickname. Identifies the initiators to which the mapping applies:
 - *initiator-name*—The mapping applies to this initiator only.
 - *initiator-ID*—The mapping applies to this initiator only, and the initiator has no nickname.
 - *host-name.**—The mapping applies to all initiators in this host.
 - *host-group-name.*.**—The mapping applies to all hosts in this group.
- Volume. Identifies the volumes to which the mapping applies:
 - *volume-name*—The mapping applies to this volume only.
 - *volume-group-name.**—The mapping applies to all volumes in this volume group.
- Access. Shows the type of access assigned to the mapping:
 - *read-write*—The mapping permits read and write access.
 - *read-only*—The mapping permits read access.
 - *no-access*—The mapping prevents access.
- LUN. Shows whether the mapping uses a single LUN or a range of LUNs (indicated by *).
- Ports. Lists the controller host ports to which the mapping applies. Each number represents corresponding ports on both controllers.

To display more information about a mapping, see [“Viewing map details” \(page 114\)](#).

Creating an initiator

You can manually create initiators. For example, you might want to define an initiator before a controller port is physically connected through a switch to a host.

To create an initiator

1. Determine the FC or SAS WWN or iSCSI IQN to use for the initiator.
2. In the Hosts topic, select **Action > Create Initiator**. The Create Initiator panel opens.
3. In the Initiator ID field, enter the WWN or IQN. A WWN value can include a colon between each pair of digits but the colons will be discarded.
4. In the Initiator Name field, enter a nickname that helps you easily identify the initiator. For example, you could use `MailServer_FCp1`. An initiator name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system or include the following: " , . < \
- If the name is used by another initiator, you are prompted to enter a different name.
5. In the Profile list, select the appropriate option:
 - o **Standard**. Default profile.
 - o **HP-UX**. The host uses Flat Space Addressing.
 - o **OpenVMS**. LUN 0 cannot be assigned to an initiator.
6. Click **OK**. The initiator is created and the hosts table is updated.

Modifying an initiator

You can modify manually created initiators.

To modify an initiator

1. In the Hosts topic, select one initiator to modify.
2. Select **Action > Modify Initiator**. The Modify Initiator panel opens.
3. In the Initiator Name field, enter a new nickname to help you identify the initiator. For example, you could use `MailServer_FCp2`. An initiator name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system or include the following: " , . < \
- If the name is used by another initiator, you are prompted to enter a different name.
4. In the Profile list, select the appropriate option:
 - o **Standard**. Default profile.
 - o **HP-UX**. The host uses Flat Space Addressing.
 - o **OpenVMS**. LUN 0 cannot be assigned to an initiator.
5. Click **OK**. The hosts table is updated.

Deleting initiators

You can delete manually created initiators that are not grouped or are not mapped. You cannot delete manually created initiators that are mapped. You also cannot delete a discovered initiator but you can remove its nickname through the delete operation.

To delete initiators

1. In the Hosts topic, select 1–1024 ungrouped, undiscovered initiators to delete.
2. Select **Action > Delete Initiators**. The Delete Initiators panel opens and lists the initiators to be deleted.
3. Click **OK**.

4. If the initiator you are trying to delete is currently undiscovered, the changes are processed and the hosts table is updated.
5. If the initiator you are trying to delete is currently discovered then a confirmation panel appears. Click **Yes** to save your changes. Otherwise, click **No**. The changes are processed and the hosts table is updated.

Adding initiators to a host

You can add existing named initiators to an existing host or to a new host.

To add an initiator to a host, the initiator must be mapped with the same access, port, and LUN settings to the same volumes or volume groups as every other initiator in the host.

To add initiators to a host

1. In the Hosts topic, select 1–128 named initiators to add to a host.
2. Select **Action > Add to Host**. The Add to Host panel opens.
3. Perform one of the following:
 - o To use an existing host, select its name in the Host Select list.
 - o To create a host, enter a name for the host in the Host Select field. A host name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system or include the following: " , . < \
4. Click **OK**. For the selected initiators, the Host value changes from -- to the specified host name.

Removing initiators from hosts

You can remove all except the last initiator from a host. Removing an initiator from a host will ungroup the initiator but will not delete it. To remove all initiators, remove the host.

To remove initiators from hosts

1. In the Hosts topic, select 1–1024 initiators to remove from their hosts.
2. Select **Action > Remove from Host**. The Remove from Host panel opens and lists the initiators to be removed.
3. Click **OK**. For the selected initiators, the Host value changes to --.

Removing hosts

You can remove hosts that are not grouped. Removing a host will ungroup its initiators but will not delete them.

To remove hosts

1. In the Hosts topic, select 1–512 ungrouped hosts to remove.
2. Select **Action > Remove Host**. The Remove Host panel opens and lists the hosts to be removed.
3. Click **OK**. For initiators that were in the selected hosts, the Host value changes to --.

Renaming a host

You can rename a host.

To rename a host

1. In the Hosts topic, select an initiator that belongs to the host that you want to rename.
2. Select **Action > Rename Host**. The Rename Host panel opens.
3. In the New Host Name field, enter a new name for the host. A host name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system or include the following: " , . < \
- If the name is used by another host, you are prompted to enter a different name.
4. Click **OK**. The hosts table is updated.

Adding hosts to a host group

You can add existing hosts to an existing host group or new host group.

To add a host to a host group, the host must be mapped with the same access, port, and LUN settings to the same volumes or volume groups as every other initiator in the host group. This means that the host must be mapped with the same access, port, and LUN settings to the same volumes or volume groups.

To add hosts to a host group

1. In the Hosts topic, select 1–256 initiators that belong to a host that you want to add to a host group.
2. Select **Action > Add to Host Group**. The Add to Host Group panel opens.
3. Perform one of the following:
 - o To use an existing host group, select its name in the Host Group Select list.
 - o To create a host group, enter a name for the host group in the Host Group Select field. A host group name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system or include the following:
", . < \
4. Click **OK**. For the selected hosts, the Group value changes from -- to the specified host group name.

Removing hosts from a host group

You can remove all except the last host from a host group. Removing a host from a host group will ungroup the host but will not delete it. To delete a host group, see [“Removing host groups” \(page 83\)](#).

To remove hosts from a host group

1. In the Hosts topic, select 1–256 hosts to remove from their host group.
2. Select **Action > Remove from Host Group**. The Remove from Host Group panel opens and lists the hosts to be removed.
3. Click **OK**. For the selected hosts, the Group value changes to --.

Renaming a host group

You can rename a host group.

To rename a host group

1. In the Hosts topic, select a host group to rename.
2. Select **Action > Rename Host Group**. The Rename Host Group panel opens.
3. In the New Host Group Name field, enter a new name for the host group. A host group name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system or include the following: ", . < \
- If the name is used by another host group, you are prompted to enter a different name.
4. Click **OK**. The hosts table is updated.

Removing host groups

You can remove host groups. Removing a host group will ungroup its hosts but will not delete them.

To remove host groups

1. In the Hosts topic, select 1–32 host groups to remove.
2. Select **Action > Remove Host Group**. The Remove Host Group panel opens and lists the host groups to be removed.
3. Click **OK**. For hosts that were in the selected host groups, the Group value changes to --.

Configuring CHAP

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 identification, 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.

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 for all iSCSI nodes, see [“Changing host port settings” \(page 56\)](#).

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. As the originator, with a valid CHAP record it can authenticate CHAP even if CHAP is disabled. This is possible because the system will supply the CHAP secret requested by its peer and the connection will be allowed. For information about setting up CHAP for use in a peer connection and how CHAP interacts with replication, see [“Creating a peer connection” \(page 123\)](#).


To add or modify a CHAP record

1. If you intend to use mutual CHAP and need to determine the IQN of a controller iSCSI port, perform the following:
 - o Select the System topic.
 - o Select the Rear view.
 - o Hover the cursor over the iSCSI host port that you intend to use. In the Port Information panel that appears, note the IQN in the ID field value.
2. In the Hosts topic, select **Action > Configure CHAP**. The Configure CHAP panel opens with existing CHAP records listed.
3. Select the **Enable Authentication (CHAP)** checkbox to enable use of CHAP for all iSCSI nodes, then confirm the operation.

NOTE: Enabling or disabling CHAP here will update its setting in the Advanced Settings tab in the Host Ports Settings panel.

4. Perform one of the following:
 - o To modify an existing record, select it. The record values appear in the fields below the CHAP records list for editing. You cannot edit the IQN.
 - o To add a new record, click **New**.
5. For a new record, in the Node Name (IQN) field, enter the IQN of the initiator. The value is case sensitive and can include a maximum of 223 bytes, including 0–9, lowercase a–z, hyphen, colon, and period.
6. In the Secret field, enter a secret for the target to use to authenticate the initiator. The secret is case sensitive and can include 12–16 bytes. The value can include spaces and printable UTF-8 characters except for the following: " <
7. To use mutual CHAP:
 - o Select the **Mutual CHAP** check box.
 - o In the Mutual CHAP Name field, enter the IQN obtained in step 1. The value is case sensitive and can include a maximum of 223 bytes and the following: 0–9, lowercase a–z, hyphen, colon, and period.
 - o In the Mutual CHAP Secret field, enter a secret for the initiator to use to authenticate the target. The secret is case sensitive, can include 12–16 bytes, and must differ from the initiator secret. The value can include spaces and printable UTF-8 characters except for the following: " <
A storage system secret is shared by both controllers.
8. Click **Apply** or **OK**. The CHAP records table is updated.

To delete a CHAP record

 **CAUTION:** Deleting CHAP records may make volumes inaccessible and the data in those volumes unavailable.

1. In the Hosts topic, select **Action > Configure CHAP**. The Configure CHAP panel opens with existing CHAP records listed.
2. Select the record to delete.
3. Click **Delete**. A confirmation panel appears.
4. Click **Remove** to continue. Otherwise, click **Cancel**. If you clicked Remove, the CHAP record is deleted.

5 Working in the Pools topic






Viewing pools

The Pools topic shows a tabular view of information about the pools and disk groups that are defined in the system, as well as information for the disks that each disk group contains. There is another type of disk group, the read-cache disk group, which is also related to virtual storage. Read-cache disk groups consist of SSDs. If your system does not use SSDs, you will not be able to create read-cache disk groups.

For information about using tables, see [“Tips for using tables” \(page 12\)](#). For more information about pools, see [“About pools” \(page 22\)](#). For more information about disk groups, see [“About disk groups” \(page 17\)](#).

Pools table

The pools table shows the following information. The system is limited to two virtual pools, which are named A and B.

- Name. Shows the name of the pool.
- Health. Shows the health of the pool:  OK,  Degraded,  Fault,  N/A, or  Unknown.
- Total Size. Shows the storage capacity defined for the pool when it was created.
- Avail. Shows the storage capacity presently available for the pool.
- Volumes. Shows the number of volumes defined for the disk groups of the pool.
- Disk Groups. Shows the number of disk groups that the pool has.

To see more information about a pool, hover the cursor over the pool in the table. The Pool Information panel that appears contains the following information:






Pool Information	Name, serial number, size, available, overcommit, pool overcommitted, low threshold, mid threshold, high threshold, allocated pages, snapshot pages, available pages, sector format, health
------------------	---

For more information about and to manage the above overcommit, low threshold, mid threshold, and high threshold settings, see [“Changing pool settings” \(page 93\)](#).

Related Disk Groups table

When you select a pool in the pools table, the disk groups for it appear in the Related Disk Groups table.

For selected pools, the Related Disk Groups table shows the following information.

- Name. Shows the name of the disk group.
- Health. Shows the health of the disk group:  OK,  Degraded,  Fault,  N/A, or  Unknown.
- Pool. Shows the name of the pool to which the disk group belongs.
- RAID. Shows the RAID level for the disk group.
- Disk Type. Shows the disk type. For virtual disk groups, the disk group's tier appears in parentheses after its disk type. For read-cache disk groups, Read Cache appears in parentheses after the disk type.
- Size. Shows the storage capacity defined for the disk group when it was created.
- Free. Shows the available storage capacity for the disk group.
- Current Job. Shows the following current system operations for the disk group, if any are occurring:
 - DRSC: Disks in the disk group are being scrubbed.
 - INIT: The disk group is being initialized.
 - RCON: At least one disk in the disk group is being reconstructed.
 - VDRAIN: The disk group is being removed and its data is being drained to another disk group.
 - VPREP: The virtual disk group is being prepared for use in a virtual pool.
 - VRECV: The virtual disk group is being recovered to restore its membership in the virtual pool.






- VREMV: The virtual disk group and its data are being removed.
- VRFY: The disk group is being verified.
- VRSC: The disk group is being scrubbed.
- Status. Shows the status for the disk group:
 - CRIT: Critical. The disk group is online but isn't fault tolerant because some of its disks are down.
 - DMGD: Damaged. The disk group is online and fault tolerant, but some of its disks are damaged.
 - FTDN: Fault tolerant with a down disk. The disk group is online and fault tolerant, but some of its disks are down.
 - FTOL: Fault tolerant and online. The disk group is online and fault tolerant.
 - MSNG: Missing. The disk group is online and fault tolerant, but some of its disks are missing.
 - OFFL: Offline. Either the disk group is using offline initialization, or its disks are down and data may be lost.
 - QTCR: Quarantined critical. The disk group is critical with at least one inaccessible disk. For example, two disks are inaccessible in a RAID-6 disk group or one disk is inaccessible for other fault-tolerant RAID levels. If the inaccessible disks come online or if after 60 seconds from being quarantined the disk group is QTCR or QTDN, the disk group is automatically dequarantined.
 - QTDN: Quarantined with a down disk. For example, the RAID-6 disk group has one inaccessible disk. The disk group is fault tolerant but degraded. If the inaccessible disks come online or if after 60 seconds from being quarantined the disk group is QTCR or QTDN, the disk group is automatically dequarantined.
 - QTOF: Quarantined offline. The disk group is offline with multiple inaccessible disks causing user data to be incomplete.
 - QTUN: Quarantined unsupported. The disk group contains data in a format that is not supported by this system. For example, this system does not support linear disk groups. For more information, see [“About the handling of linear storage” \(page 16\)](#).
 - STOP: The disk group is stopped.
 - UNKN: Unknown.
 - UP: Up. The disk group is online and does not have fault-tolerant attributes.
- Disks. Shows the number of disks in the disk group.

To see more information about a disk group, select the pool for the disk group in the pools table, then hover the cursor over the disk group in the Related Disk Groups table:

Disk Group Information	Virtual: Name, serial number, pool, tier, % of pool, allocated pages, available pages, chunk size, sector format, creation date, minimum disk size, active drive spin down enable, size, free, RAID, disks, status, current job, health
	Read cache: Name, serial number, pool, tier, allocated pages, available pages, sector format, health

Related Disks table

When you select a disk group in the Related Disk Groups table, the disks for it appear in the Related Disks table.

- Location. Shows the location of the disk.
- Health. Shows the health of the disk:  OK,  Degraded,  Fault,  N/A, or  Unknown.
- Description. Shows the disk type:
 - SAS: Enterprise SAS spinning disk.
 - SAS MDL: Midline SAS spinning disk.
 - SSD SAS: SAS solid-state disk.
- Size. Shows the storage capacity of the disk.

- Usage. Shows how the disk is being used:
 - VIRTUAL POOL: The disk is part of a virtual pool.
 - LEFTOVR: The disk is leftover.
 - FAILED: The disk is unusable and must be replaced. Reasons for this status include: excessive media errors, SMART error, disk hardware failure, or unsupported disk.
- Disk Group. Shows the disk group that contains the disk.
- Status. Shows the status of the disk:
 - Up: The disk is present and is properly communicating with the expander.
 - Spun Down: The disk is present and has been spun down by the DSD feature.
 - Warning: The disk is present but the system is having communication problems with the disk LED processor. For disk and midplane types where this processor also controls power to the disk, power-on failure will result in Error status.
 - Unrecoverable: The disk is present but has unrecoverable errors.

To see more information about a disk in a disk group, select the pool for the disk group in the pools table, select the disk group in the Related Disk Groups table, and then hover the cursor over the disk in the Related Disks table:

Disk Information	Location, serial number, usage, type, size, status, revolutions per minute (spinning disk only), SSD life left, manufacturer, model, firmware revision, power on hours, job status, FDE state (MSA 2050 only), FDE lock key (MSA 2050 only), job running, sector format, health
------------------	---

NOTE: Following is more information for selected Disk Information panel items:

- *Power On Hours* refers to the total number of hours that the disk has been powered on since it was manufactured. This value is updated in 30-minute increments.
- For MSA 2050:
FDE State refers to the FDE state of the disk. For more information about FDE states, see the CLI Reference Guide.
- For MSA 2050:
FDE lock keys are generated from the FDE passphrase and manage locking and unlocking the FDE-capable disks in the system. Clearing the lock keys and power cycling the system denies access to data on the disks.

Adding a disk group

You can create virtual disk groups using specified disks through the Add Disk Group panel. You can also create read-cache disk groups through this panel. When creating a disk group, you explicitly select the RAID level and individual disks and incorporate them into a pool. All disks in a disk group must be the same type (enterprise SAS, for example). 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). The Performance Tier license is required to create a virtual disk group comprised of SSDs for use as a Performance tier. The Performance Tier license is not required in order to use SSDs in read-cache disk groups, or in an all-flash array. An all-flash array does not use tiering. For more information about disk groups, see [“About disk groups” \(page 17\)](#).

Add Disk Group panel overview

The Add Disk Group panel is dynamic, displaying options based on the type of disk group you want to create and the data protection level selected. There are three sections that comprise the panel.

The top section provides options to name and define the disk group type, select the pool it will reside on, and choose its data protection (RAID) level.

The middle section contains the Disk Selection Sets summary which presents cumulative data for the disks selected for the disk group. It displays information about the data protection and disk type selected for the disk group, as well as the total number of disks selected, the minimum and maximum number of disks allowed for the specified data protection

level, the size of the disk group (total capacity of all selected drives), and the Complete check box. The Complete check box indicates if the minimum number of disks needed to configure the disk group have been selected, and automatically changes from to .

As you select drives to add to the disk group, a color-coded bar graph displays the disk group's available capacity, dedicated overhead capacity (for data protection and array metadata), and wasted capacity.

The bottom section lists the disks located within each enclosure in your system, along with their details. Add disks to the disk group by doing one of the following:

- Select a range of disks within an enclosure by entering a comma-separated list that contains the enclosure number and disk range in the Enter Range of Disks text box. Use the format *enclosure-number.disk-range, enclosure-number.disk-range*. For example, to select disks 3-12 in enclosure 1 and 5-23 in enclosure 2, enter **1.3-12,2.5-23**.
- Select all disks by checking the Select All checkbox.
- Filter the disks in the list per disk type, enclosure ID, slot location, disk size, or health by entering applicable search criteria in the text box. Clear the filter by selecting the Clear Filters button.
- Click on individual disks within the table to select them and add them to the disk group.

Selected disks are highlighted in blue. Remove disks from the group by clicking on them to deselect them.

Virtual disk groups

The system supports a maximum of two pools, one per controller module: A and B. You can add up to 16 virtual disk groups for each virtual pool. If a virtual pool does not exist, the system will automatically add it when creating the disk group. Once a virtual pool and disk group exist, volumes can be added to the pool. Once you add a virtual disk group, you cannot add or remove individual disks from the group. If your organization's needs change, you can modify your storage amount by adding new virtual disk groups or removing existing ones.

Depending on the type of disks selected and license installed, virtual disk groups belong to one of the following tiers:

- Enterprise SAS disks: Standard tier.
- Midline SAS disks: Archive tier.
- SSDs: Performance tier. Requires the Performance Tier license to be used as a capacity tier (non-Read Cache) and HDDs are used in the same system. Does not require the license to be used in read-cache disk groups, or in all-flash arrays, which do not use tiers.

TIP:

- All virtual disk groups in the same tier within a virtual pool should have the same configuration of number of disks, capacity, and data protection (RAID) level. This will provide consistent performance across the tier.

NOTE: If a virtual pool contains a single virtual disk group, and it has been quarantined, you cannot add a new virtual disk group to the pool until you've resolved the quarantined disk group. Contact technical support for assistance.

Read-cache disk groups

If your system has SSDs, you can also add read-cache disk groups. Read cache is a special type of virtual disk group that can be added only to a virtual pool. It is used for the purpose of caching virtual pages for improving read performance. A virtual pool can contain only one read-cache disk group. A virtual pool cannot contain both read cache and a Performance tier. At least one virtual disk group must exist before a read-cache disk group can be added. NRAID is automatically used for a read-cache disk group with a single disk. RAID-0 is automatically used for a read-cache disk group with the maximum of two disks. When you create a read-cache disk group, the system automatically creates a read-cache tier, if one does not already exist. Unlike the other tiers, it is not used in tiered migration of data.

Disk group options

The following options appear in the top section of the Add Disk Group panel:

- **Type.** When creating a disk group, select **Virtual** to show options for a virtual disk group, or **Read Cache** to show options for a read cache disk group.
- **Name.** A disk group name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system or include the following: " , < \

By default, disk groups are named `dgcontrollerxx`, where controller is either A or B and xx starts at 01, and read-cache disk groups are named `rccontrollerx`, where x is 1 or 2.

- **RAID Level.** Select one of the following RAID levels when creating a virtual disk group (the default setting is RAID 6):
 - **RAID 1.** Requires 2 disks.
 - **RAID 5.** Requires 3-16 disks.
 - **RAID 6.** Requires 4-16 disks.
 - **RAID 10.** Requires 4-16 disks, with a minimum of two RAID-1 subgroups, each having two disks.

NOTE: For a virtual group, the system will use one of the following chunk sizes, which cannot be changed:

- RAID 1: Not applicable
- RAID 5 and RAID 6:
 - With 2, 4, or 8 non-parity disks: 512k. For example, a RAID-5 group with 3, 5, or 9 total disks or a RAID-6 group with 4, 6, or 10 total disks.
 - Other configurations: 64k
- RAID 10: 512k

-
- **Pool** (only appears for virtual and read-cache disk groups). Select the name of the virtual pool (A or B) to contain the group.
 - **Number of Sub-groups** (options only appear when RAID-10 is selected). Changes the number of sub-groups that the disk group should contain.

To add a disk group

1. In the Pools topic, select **Action > Add Disk Group**. The Add Disk Group panel opens.
2. Set the options.
3. Select the disks.

NOTE: Depending on the licensing for your system and the type of disks that it contains, some or all disks might be grayed in the user interface and unavailable.

4. Click **Add**. If your disk group contains a mix of 512n and 512e disks, a dialog box displays. Perform one of the following:
 - To create the disk group, click **Yes**.
 - To cancel the request, click **No**.

If the task succeeds, the new disk group appears in the Related Disk Groups table in the Pools topic when you select the pool for it in the pools table.

Modifying a disk group

You can rename any virtual and read-cache disk group.

To modify a disk group

1. In the Pools topic, select the pool for the disk group that you are modifying in the pools table. Then, select the disk group in the Related Disk Groups table.

NOTE: To see more information about a pool, hover the cursor over the pool in the table. [“Viewing pools” \(page 86\)](#) contains more details about the Pool Information panel that appears.

2. Select **Action > Modify Disk Group**. The Modify Disk Group panel opens.
3. To change the disk group name, replace the existing name in the New Name field. A disk group name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system or include the following: " , < \
4. Click **Modify**.
5. To close the confirmation panel, click **OK**.

Removing disk groups

You can delete a single disk group or select multiple disk groups and delete them in a single operation. By removing disk groups, you can also remove pools. Removing all disk groups within a pool will also trigger the automatic removal of the associated pool.

If all disk groups for a pool have volumes assigned and are selected for removal, a confirmation panel will warn the user that the pool and all its volumes will be removed.

Unless a virtual pool consists exclusively of SSDs, if a virtual pool has more than one disk group and at least one volume that contains data, the system attempts to drain the disk group to be deleted by moving the volume data that it contains to other disk groups in the pool. When removing one or more, but not all, disk groups from a virtual pool, the following possible results can occur:

- If the other disk groups do not have room for the data of the selected disk group, the delete operation will fail immediately and a message will be displayed.
- If there is room to drain the volume data to other disk groups, a message will appear that draining has commenced and an event will be generated upon completion (progress will also be shown in the Current Job column of the Related Disk Groups table).
 - When the disk group draining completes, an event will be generated, the disk group disappears, and the drives for it becomes available.
 - If a host writes during the disk group draining, which results in there not being enough room to finish the draining, an event will be generated, the draining terminates, and the disk group will remain in the 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: large pool configuration; the amount of I/O traffic to the system (e.g., active I/O pages to the draining disk group); the type of the disk group page migration (enterprise SAS, midline SAS, SSD); the size of the draining disk group(s) in the system; and the number of disk groups draining at the same time.

If you remove the last disk group in a virtual pool, the system will prompt you to confirm removing the pool, too. If you choose yes, the pool will be removed. If you choose no, the disk group and the pool will remain.

NOTE: If the disk group is the last disk group for a pool that is used in a peer connection or it contains a volume that is used in a replication set, the **Remove Disk Groups** menu option will be unavailable.

To remove a disk group

1. In the Pools topic, select the pool for the disk group(s) that you are deleting in the pools table. Then, select the disk group(s) in the Related Disk Groups table.

NOTE: To see more information about a pool, hover the cursor over the pool in the table. [“Viewing pools” \(page 86\)](#) contains more details about the Pool Information panel that appears.

2. Select **Action > Remove Disk Groups**. The Remove Disk Groups panel opens.
3. Click **OK**. A confirmation panel displays.
4. Click **Yes** to remove all disk groups from the pool. Otherwise, click **Cancel**. If you clicked Yes, the disk group(s) and their volumes are deleted, the pool for the disk group(s) might be deleted, the disks for the disk group(s) become available, and the Related Disk Groups table is updated.
5. A message about volume migration appears. Click **OK** to confirm.

Managing Spares

The Manage Spares panel displays a list of current spares and lets you add and remove global spares.

Global spares

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 spare must have sufficient capacity to replace the smallest disk in an existing disk group.

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

The Manage Spares panel consists of two sections. The top section lists the current spares in the system and includes information about each. The bottom section lists available disks located within each enclosure in your system that can be designated as global spares and includes details about each. Click on individual disks within the table to select them. Search for specific disks in the list by entering search criteria in the text box. Filter the disks in the list per disk type, location, or disk size by entering applicable search criteria in the text box. Clear the filter by selecting the Clear Filters button.

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 chooses the spare after a disk in the disk group fails. For more information about disk groups, see [“About disk groups” \(page 17\)](#).

To add global spares

1. In the Pools topic, select **Action > Manage Spares**. The Manage Spares panel displays.
2. In the Add New Spares section, click on available disks to select them.
3. Click **Add Spares**. The system updates the global spares and a confirmation panel displays.
4. To close the confirmation panel, click **OK**.

To remove global spares

1. In the Pools topic, select **Action > Manage Spares**. The Manage Spares panel displays.
2. In the Current Spares section, click on available the spares to remove.
3. Click **Remove**. The system updates the global spares and a confirmation panel displays.
4. To close the confirmation panel, click **OK**.

Creating a volume

You can add volumes to virtual pools. The Create Virtual Volumes panel enables you to create virtual volumes. You can access these panels from both the Pools and Volumes topics.

To create volumes through the Pools topic

1. In the Pools topic, select a pool in the pools table.

NOTE: To see more information about a pool, hover the cursor over the pool in the table. [“Viewing pools” \(page 86\)](#) contains more details about the Pool Information panel that appears.

2. Select **Action > Create Volumes**. The Create Virtual Volumes panel opens.
3. For more information about creating virtual volumes, see [“Creating a virtual volume” \(page 99\)](#).

Changing pool settings

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

NOTE: If the pool size is 500 GB or smaller, and/or the middle threshold is relatively high, the high threshold may not guarantee 200 GB of reserved space in the pool. The controller will not automatically adjust the low and middle thresholds in such cases.

You can view and change settings that govern the operation of each virtual pool:

- **Low Threshold.** When this percentage of virtual pool capacity has been used, informational event 462 will be generated to notify the administrator. This value must be less than the Mid Threshold value. The default is 50%.
- **Mid Threshold.** When this percentage of virtual pool capacity has been used, event 462 will be generated 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 event will have Informational severity. If the pool is overcommitted, the event will have Warning severity.
- **High Threshold.** When this percentage of virtual pool capacity has been used, event 462 will be generated to alert the administrator to add capacity to the pool. This value is automatically calculated based on the available capacity of the pool minus 200 GB of reserved space. If the pool is not overcommitted, the event will have Informational severity. If the pool is overcommitted, the event will have Warning severity and the system will use write-through cache mode until virtual pool usage drops back below this threshold.
- **Enable overcommitment of pools?.** This check box controls whether thin provisioning is enabled, and whether storage-pool capacity may exceed the physical capacity of disks in the system. For information about thin provisioning, see [“About thin provisioning” \(page 24\)](#). This option is enabled by default.

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. If the pool is overcommitted and has exceeded its high threshold, its health will show as degraded in the Pools topic. 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, in the Pools topic, display the Pool Information panel by hovering the cursor over the pool in the pools table. In that panel, if the Pool Overcommitted value is `True`, the pool is overcommitted. If the value is `False`, the pool is not overcommitted.

To change virtual pool settings

1. In the Pools topic, select a virtual pool in the pools table.

NOTE: To see more information about a virtual pool, hover the cursor over the pool in the table. [“Viewing pools” \(page 86\)](#) contains more details about the Pool Information panel that appears.

2. Select **Action > Change Pool Settings**. The Pool Settings panel opens.
3. To change the low and mid thresholds for each pool, enter new values.
4. To enable thin provisioning, select the **Enable overcommitment of pool?** check box.
5. Click **OK**. The changes are saved.

Verifying and scrubbing disk groups

Verifying a disk group

If you suspect that a fault-tolerant (mirror or parity) disk group has a problem, run the Verify utility to check the disk group's integrity. For example, if you haven't checked the system for parity inconsistencies recently and are concerned about the disk health, verify its disk groups. The Verify utility analyzes the selected disk group to find and fix inconsistencies between its redundancy data and its user data. This utility fixes parity mismatches for RAID 5, 6, and find but not fix mirror mismatches for RAID 1 and 10. This task can be performed only on a disk group whose status is FTOL (fault tolerant and online). It cannot be performed for NRAID or RAID 0 read cache disk groups.

Verification can last over an hour, depending on the size of the disk group, the utility priority, and the amount of I/O activity. You can use a disk group while it is being verified. When verification is complete, event 21 is logged and specifies the number of inconsistencies found. Such inconsistencies can indicate that a disk in the disk group is going bad. For information about identifying a failing disk, use the SMART option. For more information, see [“Configuring SMART” \(page 73\)](#).

If too many utilities are running for verification to start, either wait until those utilities have completed and try again, or abort a utility to free system resources. If you abort verification, you cannot resume it. You must start it over.

To verify a disk group

1. In the Pools topic, select the pool for the disk group that you are verifying in the pools table. Then, select the disk group in the Related Disk Groups table.

NOTE: To see more information about a pool, hover the cursor over the pool in the table. [“Viewing pools” \(page 86\)](#) contains more details about the Pool Information panel that appears.

2. Select **Action > Disk Group Utilities**. The Disk Group Utilities panel opens, showing the current job status.
3. Click **Verify Disk Group**. A message confirms that verification has started.
4. Click **OK**. The panel shows the verification's progress.

To abort disk group verification

1. In the Pools topic, select the pool for the disk group that you are verifying in the pools table. Then, select the disk group in the Related Disk Groups table.
2. Select **Action > Disk Group Utilities**. The Disk Group Utilities panel opens, showing the current job status.
3. Click **Abort Verify**. A message confirms that verification has been aborted.
4. Click **OK**.

Scrubbing a disk group

The system-level Disk Group Scrub option automatically checks all disk groups for disk defects. If this option is disabled, you can still perform a scrub on a selected disk group. Scrub analyzes the selected disk group to find and fix disk errors. It will fix parity mismatches for RAID 5 and 6 and mirror mismatches for RAID 1 and 10.

Scrub can last over an hour, depending on the size of the disk group, the utility priority, and the amount of I/O activity. However, a manual scrub performed by Scrub Disk Group is typically faster than a background scrub performed by Disk Group Scrub. You can use a disk group while it is being scrubbed. When a scrub is complete, event 207 is logged and specifies whether errors were found and whether user action is required.

To scrub a disk group

1. In the Pools topic, select the pool for the disk group that you are verifying in the pools table. Then, select the disk group in the Related Disk Groups table.
2. Select **Action > Disk Group Utilities**. The Disk Group Utilities panel opens, showing the current job status.
3. Click **Scrub Disk Group**. A message confirms that the scrub has started.
4. Click **OK**. The panel shows the scrub's progress.

To abort a disk group scrub

1. In the Pools topic, select the pool for the disk group that you are verifying in the pools table. Then, select the disk group in the Related Disk Groups table.

NOTE: If the disk group is being scrubbed but the Abort Scrub button is grayed out, a background scrub is in progress. To stop the background scrub, disable the Disk Group Scrub option as described in [“Configuring system utilities” \(page 77\)](#).

2. Select **Action > Disk Group Utilities**. The Disk Group Utilities panel opens, showing the current job status.
3. Click **Abort Scrub**. A message confirms that the scrub has been aborted.
4. Click **OK**.

6 Working in the Volumes topic

Viewing volumes

The Volumes topic shows a tabular view of information about volumes, replication sets, and snapshots that are defined in the system. For more information about volumes, see [“About volumes and volume groups” \(page 23\)](#). For more information about replication, see [“About replicating virtual volumes” \(page 116\)](#). For more information about snapshots, see [“About snapshots” \(page 27\)](#). For information about using tables, see [“Tips for using tables” \(page 12\)](#).

Volumes table


To see more information about a volume or snapshot, hover the cursor over an item in the volumes table. The Volume Information panel opens with more detailed information about the item. The following table displays the categories of information while descriptions for selected terms follow.

Volume Information	Name, type, pool, group, class, size, allocated size, owner, serial number, volume copy job, write policy, optimization, read-ahead size, tier affinity, health
--------------------	---

For more information about write policy and read-ahead size, see [“Modifying a volume” \(page 100\)](#).

The volumes table shows the following information. By default, the table shows 10 entries at a time.

- Group. Shows the group name if the volume is grouped into a volume group; otherwise, --.
- Name. Shows the name of the volume.
- Pool. Shows whether the volume is in pool A or B for virtual pools.
- Type. Shows whether the volume is a base volume (virtual) or a snapshot (virtual).
- Size. Shows the storage capacity defined for the volume when it was created (minus 60 KB for internal use).
- Allocated. Shows the storage capacity allocated to the volume for written data.

 **TIP:** When selecting one or more volumes or snapshots in the volumes table, the **Snapshots**, **Maps**, **Replication Sets**, and **Schedules** tabs will be enabled if they have associated information for the selected items. They will be grey and disabled if they do not.

Snapshots table

To see more information about a snapshot and any child snapshots taken of it, select the snapshot or volume that is associated with it in the volumes table. If it is not already selected, select the **Snapshots** tab. The snapshots and all related snapshots appear in the Snapshots table. Then, hover the cursor over the item in the Snapshots table:

Snapshot Information	Virtual: Name, serial number, status, status reason, retention priority, snapshot data, unique data, shared data, pool, class, number of snaps, number of snapshots in tree, source volume, total size, creation date/time, type, parent volume, base volume, health
----------------------	--

The Snapshots table shows the following snapshot information. By default, the table shows 10 entries at a time.

- Name. Shows the name of the snapshot.
- Base Volume. Shows the name of the virtual volume from which the snapshot was created. All virtual volumes are base volumes when created and are volumes from which virtual snapshots can be created.
- Parent Volume. Shows the name of the volume from which the snapshot was created.
- Creation Date/Time. Shows the date and time when the snapshot was created.

- **Status.** Shows whether the snapshot is available or unavailable. A snapshot can be unavailable for one of the following reasons:
 - The source volume is not accessible or is not found.
 - The snapshot is pending.
 - A rollback with modified data is in progress.
- **Snap Data.** Shows the total amount of data associated with the specific snapshot (data copied from a source volume to a snapshot and data written directly to a snapshot).

Maps table

To see information about the maps for a snapshot or volume, select the snapshot or volume in the volumes table. Then, select the **Map** tab. The maps appear in the Maps table.

The Maps table shows the following mapping information. By default, the table shows 10 entries at a time.

- **Group.Host.Nickname.** Identifies the initiators to which the mapping applies:
 - *initiator-name*—The mapping applies to this initiator only.
 - *initiator-ID*—The mapping applies to this initiator only, and the initiator has no nickname.
 - *host-name.**—The mapping applies to all initiators in this host.
 - *host-group-name.*.**—The mapping applies to all hosts in this group.
- **Volume.** Identifies the volumes to which the mapping applies:
 - *volume-name*—The mapping applies to this volume only.
 - *volume-group-name.**—The mapping applies to all volumes in this volume group.
- **Access.** Shows the type of access assigned to the mapping:
 - *read-write*—The mapping permits read and write access.
 - *read-only*—The mapping permits read access.
 - *no-access*—The mapping prevents access.
- **LUN.** Shows the LUN number or ' * ' if the map is to a volume group.
- **Ports.** Lists the controller host ports to which the mapping applies. Each number represents corresponding ports on both controllers.


To display more information about a mapping, see [“Viewing map details” \(page 114\)](#).


Replication Sets table

To see information about the replication set for a volume or volume group, select a volume in the volumes table. If it is not already selected, select the **Replication Sets** tab. The replication appears in the Replication Sets table. To see more information about the replication set, hover the cursor over each item in the table:

Replication Set Information	Name, serial number, status, primary volume group, primary volume group serial, secondary volume group, secondary volume group serial, peer connection, queue policy, queue count, secondary volume snapshot history, primary volume snapshot history, retention count, retention priority, snapshot basename, associated schedule name, current run progress, current run start time, current run estimated time to completion, current run transferred data, last successful run, last run start time, last run end time, last run transferred data, last run status, last run error status
-----------------------------	---

The Replication Sets table shows the following information. By default, the table shows 10 entries at a time.

- **Name.** Shows the replication set name.
- **Primary Volume.** Shows the primary volume name. For replication sets that use volume groups, the primary volume name is *volume-group-name.** where *.** signifies that the replication set contains more than one volume. If the volume is on the local system, the  icon appears.

- **Secondary Volume.** Shows the secondary volume name. For replication sets that use volume groups, the secondary volume name is `volume-group-name.*` where `.` signifies that the replication set contains more than one volume. If the volume is on the local system, the  icon appears.
- **Status.** Shows the status of the replication set:
 - **Not Ready.** The replication set is not ready for replications because the system is still preparing the replication set.
 - **Unsynchronized.** The primary and secondary volumes are unsynchronized because the system has prepared the replication set, but the initial replication has not run.
 - **Running.** A replication is in progress.
 - **Ready.** The replication set is ready for a replication.
 - **Suspended.** Replications have been suspended.
 - **Unknown:** This system cannot communicate with the primary system and thus cannot be sure of the current state of the replication set. Check the state of the primary system.
- **Last Successful Run.** Shows the date and time of the last successful replication.
- **Estimated Completion Time.** Shows the estimated date and time for the replication in progress to complete.

Schedules table

For information about the schedules for a snapshot, select the snapshot in the volumes table. For information about the schedules for copy operations for a volume, select the volume in the volumes table. For information about the schedules for a replication set, select a volume for the replication set in the volumes table. If it is not already selected, select the **Schedules** tab. The schedules appear in the Schedules table. Then, hover the cursor over the item in the Schedules table.

Schedule Information	Name, schedule specification, schedule status, next time, task name, task type, task status, task state, error message. Additional schedule information per task type:
	<ul style="list-style-type: none"> ◦ Replication set - source volume, source volume serial ◦ Reset snapshot - snapshot name, snapshot serial ◦ Take snapshot - source volume, source volume serial, prefix, count, last created

The Schedules table shows the following schedule information. By default, the table shows 10 entries at a time.

- **Schedule Name.** Shows the name of the schedule.
- **Schedule Specification.** Shows the schedule settings for running the associated task.
- **Status.** Shows the status for the schedule:
 - **Uninitialized.** The schedule is not yet ready to run.
 - **Ready.** The schedule is ready to run at the next scheduled time.
 - **Suspended.** The schedule had an error and is holding in its current state.
 - **Expired.** The schedule exceeded a constraint and will not run again.
 - **Invalid.** The schedule is invalid.
 - **Deleted.** The schedule has been deleted.
- **Task Type.** Shows the type of schedule:
 - **TakeSnapshot.** The schedule creates a snapshot of a source volume.
 - **ResetSnapshot.** The schedule deletes the data in the snapshot and resets it to the current data in the volume from which the snapshot was created. The snapshot's name and other volume characteristics are not changed.
 - **VolumeCopy.** The schedule copies a source volume to a new volume. It creates the destination volume you specify, which must be in a disk group owned by the same controller as the source volume. The source volume can be a base volume or a snapshot.
 - **Replicate.** The schedule replicates a virtual replication set to a remote system.

Creating a virtual volume

You can add volumes to a virtual pool. You can create an individual virtual volume, multiple virtual volumes with different settings, or multiple virtual volumes with the same settings. In the latter case, the volumes will have the same base name with a numeric suffix (starting at 0000) to make each name unique and they will be placed in the same pool. You can also select a volume tier affinity setting to specify a tier for the volume data.

The Create Virtual Volumes panel contains a graphical representation of storage capacity for pools A and B. Each graph provides the number of existing volumes, free space, allocated and unallocated space, and committed and overcommitted space for pool A or B. The graph for the specified pool of the prospective new virtual volume also shows the impact of storage space and the prospective new volume on the pool.

The volumes table in the Volumes topic lists all volumes, volume groups, and snapshots. To see more information about a virtual volume, hover the cursor over the volume in the table. [“Viewing volumes” \(page 96\)](#) contains more details about the Volume Information panel that appears.

To create virtual volumes

1. Perform one of the following:

- o In the Pools topic, select a virtual pool in the pools table and select **Action > Create Volumes**.
- o In the Volumes topic, select **Action > Create Virtual Volumes**.

The Create Virtual Volumes panel opens and shows the current capacity usage of each pool.

NOTE: If a virtual pool does not exist, the option to create virtual volumes will be unavailable.

2. Optional: Change the volume name. The default is Vol n , where n starts at 0001 and increments by one for each volume that has a default name. A volume name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system or include the following: " , < \

If the name is used by another volume, the name is automatically changed to be unique. For example, MyVolume would change to MyVolume0001, or Volume2 would change to Volume3.

3. Optional: Change the volume size, including unit of measurement. You can use any of the following units: MiB, GiB, TiB, MB, GB, TB. The default size is 100 GB. For the maximum volume size that the system supports, see the system configuration limits topic in the SMU help.

Volume sizes are aligned to 4.2-MB (4-MiB) boundaries. When a volume is created or expanded, if the resulting size is less than 4.2 MB it will be increased to 4.2 MB. A value greater than 4.2 MB will be decreased to the nearest 4.2-MB boundary.

4. Optional: Change the number of volumes to create. The default is 1. See the system configuration limits topic in SMU help for the maximum number of volumes supported per pool.

5. Optional: Specify a volume tier affinity setting to automatically associate the volume data with a specific tier, moving all volume data to that tier whenever possible. The default is **No Affinity**. For more information on the volume tier affinity feature, see [“About automated tiered storage” \(page 25\)](#).

6. Optional: Select the pool in which to create the volume. The system load-balances volumes between the pools so the default may be A or B, whichever contains fewer volumes.

7. Optional: To create another volume with different settings, click **Add Row** and then change the settings. To remove the row that the cursor is in, click **Remove Row**.

8. Click **OK**.

If creating the volume will overcommit the pool capacity, the system will prompt you to configure event notification to be warned before the pool runs out of physical storage.

9. If the virtual volume exceeds the capacity:

- a. Click **Yes** to continue. Otherwise, click **No**. If you clicked Yes, the volumes are created and the volumes table is updated.
- b. To close the confirmation panel, click **OK**.

Modifying a volume

You can change the name and cache settings for a volume. You can also expand a volume. If a virtual volume is not a secondary volume involved in replication, you can expand the size of the volume but not make it smaller. Because volume expansion does not require I/O to be stopped, the volume can continue to be used during expansion.

The volume cache settings consist of the write policy, cache optimization mode, and read-ahead size. For more information on volume cache settings, see [“About volume cache options” \(page 23\)](#).

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

The volume tier affinity settings are No Affinity, Archive, and Performance. For more information about these settings, see [“Volume tier affinity feature” \(page 25\)](#).

To see more information about a volume, hover the cursor over the volume in the table. [“Viewing volumes” \(page 96\)](#) contains more details about the Volume Information panel that appears.

To modify a volume

1. In the Volumes topic, select a volume in the volumes table.
2. Select **Action > Modify Volume**. The Modify Volume panel opens.
3. Optional: In the New Name field, enter a new name for the volume. A volume name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system or include the following: " , < \
4. Optional: In the Expand By field, enter the size by which to expand the volume. If overcommitting the physical capacity of the system is not allowed, the value cannot exceed the amount of free space in the storage pool. You can use any of the following units: MiB, GiB, TiB, MB, GB, TB.

Volume sizes are aligned to 4.2-MB (4-MiB) boundaries. When a volume is created or expanded, if the resulting size is less than 4.2 MB it will be increased to 4.2 MB. A value greater than 4.2 MB will be decreased to the nearest 4.2-MB boundary.

5. Optional: In the Write Policy list, select **Write-back** or **Write-through**. The default is **Write-back**.
6. Optional: In the Write Optimization list, select **Standard** or **No-mirror**. The default is **Standard**.
7. Optional: In the Read Ahead Size list, select **Adaptive**, **Disabled**, **Stripe**, or a specific size (512 KB; 1, 2, 4, 8, 16, or 32 MB). The default is **Default**.
8. Optional: In the Tier Affinity field, select **No Affinity**, **Archive**, or **Performance**. The default is **No Affinity**.
9. Click **OK**.

If a change to the volume size will overcommit the pool capacity, the system will prompt you to configure event notification to be warned before the pool runs out of physical storage.

10. If the virtual volume exceeds the capacity:
 - a. Click **Yes** to continue. Otherwise, click **No**. If you clicked Yes, the volumes table is updated.
 - b. To close the confirmation panel, click **OK**.

Copying a volume or snapshot

You can copy a volume or a snapshot to a new volume. To ensure the integrity of a copy, unmount the source or, at minimum, perform a system cache flush on the host and refrain from writing to the source. Since the system cache flush is not natively supported on all operating systems, it is recommended to unmount temporarily. The copy will contain all data on disk at the time of the request, so if there is data in the OS cache, that data will not be copied. Unmounting the source forces the cache flush from the host OS. After the copy has started, it is safe to remount the source and resume I/O.

To copy a virtual volume or snapshot

1. In the Volumes topic, select a virtual volume or snapshot.

2. Select **Action > Copy Volume**. The Copy Volume panel opens.
3. Optional: In the New Volume field, change the name for the new volume. The default is `volume-namecn`, where `n` starts at 01. A volume name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system or include the following: `"`, `<`, `\`
If the name is used by another volume, you are prompted to enter a different name.
4. In the Residing on Pool field, select the pool in which to create the copy. Selecting **Auto** copies the destination volume to the same pool as the source volume.
5. Click **OK**. A confirmation panel appears.
6. Click **Yes** to continue. Otherwise, click **No**.

Aborting a volume copy

You can abort a volume copy operation. When the operation is complete, the destination volume is deleted.

To abort a volume copy

1. In the Volumes topic, select a volume that is currently being copied.
2. Select **Menu > Abort Volume Copy**.
3. Click **Yes** to abort the operation.

Adding volumes to a volume group

You can add virtual volumes to a new or existing virtual volume group. All volumes in a volume group must be in the same pool.

To add a volume to a volume group, the volume must have the same mappings as all other members of the group. This means that the volume must be mapped with the same access and port settings to the same initiators, hosts, or host groups.

If the volume group is part of a replication set, you cannot add or remove volumes to or from it. If a volume group is being replicated, the maximum number of volumes that can exist in the group is 16.

NOTE: You cannot map LUN 0 for a SAS initiator. You can create a maximum of 1024 volumes, but because the supported LUN range is 1–1023, only 1023 volumes can be mapped using default mapping. Using explicit mapping, all volumes can be mapped.

To add volumes to a volume group

1. In the Volumes topic, select 1–20 volumes to add to a volume group.
2. Select **Action > Add to Volume Group**. The Add to Volume Group panel opens.
3. Perform one of the following:
 - o To use an existing volume group, select its name in the Volume Groups list.
 - o To create a volume group, enter a name for the volume group in the Volume Groups field. A volume group name is case sensitive and can have a maximum of 32 bytes. It cannot include the following: `"`, `<`, `\`
4. Click **OK**. For the selected volumes, the Volume Groups value changes from `--` to the specified host group name.

Removing volumes from a volume group

You can remove volumes from a volume group. You cannot remove all volumes from a group. At least one volume must remain. Removing a volume from a volume group will ungroup the volumes but will not delete them. To remove all volumes from a volume group, see [“Removing volume groups” \(page 102\)](#).

To see more information about a volume, hover the cursor over the volume in the table. [“Viewing volumes” \(page 96\)](#) contains more details about the Volume Information panel that appears.

To remove volumes from a volume group

1. In the Volumes topic, select the volumes to remove from a volume group.
2. Select **Action > Remove from Volume Group**. The Remove from Volume Group panel opens and lists the volumes to be removed.
3. Click **OK**. For the selected volumes, the Group value changes to --.

Renaming a volume group


You can rename a volume group unless it is part of a replication set. To see more information about a volume, hover the cursor over the volume in the table. [“Viewing volumes” \(page 96\)](#) contains more details about the Volume Information panel that appears, including how to view volumes and volume groups that are part of a replication set.

To rename a volume group

1. In the Volumes topic, select a volume that belongs to the volume group that you want to rename.
 2. Select **Action > Rename Volume Group**. The Rename Volume Group panel opens.
 3. In the New Group Name field, enter a new name for the volume group. A volume group name is case sensitive and can have a maximum of 32 bytes. It cannot include the following: " , < \
- If the name is used by another volume group, you are prompted to enter a different name.
4. Click **OK**. The volumes table is updated.

Removing volume groups

You can remove volume groups. When you remove a volume group, you can optionally delete its volumes. Otherwise, removing a volume group will ungroup its volumes but will not delete them.

 **CAUTION:** Deleting a volume removes its mappings and schedules and deletes its data.

To see more information about a volume, hover the cursor over the volume in the table. [“Viewing volumes” \(page 96\)](#) contains more details about the Volume Information panel that appears.

To remove volume groups only

1. In the Volumes topic, select a volume that belongs to each volume group that you want to remove. You can remove 1–100 volume groups at a time.
2. Select **Action > Remove Volume Group**. The Remove Volume Group panel opens and lists the volume groups to be removed.
3. Click **OK**. For volumes that were in the selected volume groups, the Volume Groups value changes to --.

To remove volume groups and their volumes

1. Verify that hosts are not accessing the volumes that you want to delete.
2. In the Volumes topic, select a volume that belongs to each volume group that you want to remove. You can remove 1–100 volume groups at a time.
3. Select **Action > Remove Volume Group**. The Remove Volume Group panel opens and lists the volume groups to be removed.
4. Select the **Delete Volumes** check box.
5. Click **OK**. A confirmation panel appears.
6. Click **Yes** to continue. Otherwise, click **No**.

If you clicked **Yes**, the volume groups and their volumes are deleted and the volumes table is updated.

Rolling back a volume

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

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

Only one rollback is allowed on the same volume at one time. Additional rollbacks are queued until the current rollback is complete. However, after the rollback is requested, the volume is available for use as if the rollback has already completed.

For volumes and snapshots, 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 the rollback. Since virtual snapshots are copies of a point in time, they 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.

You cannot roll back a volume that is part of a replication set. To see more information about a volume, hover the cursor over the volume in the table. [“Viewing volumes” \(page 96\)](#) contains more details about the Volume Information panel.

To roll back a volume

1. Unmount the volume from hosts.
2. In the Volumes topic, select the volume to roll back.
3. Select **Action > Rollback Volume**. The Rollback Volume panel opens and lists snapshots of the volume.
4. Select the snapshot to roll back to.
5. Click **OK**. A confirmation panel appears.
6. Click **Yes** to continue. Otherwise, click **No**. If you clicked **Yes**, the rollback starts. You can now remount the volume.

Deleting volumes and snapshots

You can delete volumes and snapshots. You can delete a volume that has no child snapshots. You cannot delete a virtual volume that is part of a replication set.

△ CAUTION: Deleting a volume or snapshot removes its mappings and schedules and deletes its data.

NOTE: You can only delete a volume with one or more snapshots, or a snapshot with child snapshots, by deleting all of the snapshots or child snapshots first.

To see more information about a volume or snapshot, hover the cursor over the item in the volumes table.

You can view additional snapshot information by hovering the cursor over the snapshot in the Related Snapshots table. [“Viewing volumes” \(page 96\)](#) contains more details about the Volume Information and Snapshot Information panels that appear.

To delete volumes and snapshots

1. Verify that hosts are not accessing the volumes and snapshots that you want to delete.
2. In the Volumes topic, select 1–100 items (volumes, snapshots, or both) to delete.
3. Select **Action > Delete Volumes**. The Delete Volumes panel opens with a list of the items to be deleted.
4. Click **Delete**. The items are deleted and the volumes table is updated.

Creating snapshots

You can create snapshots of selected volumes or snapshots. (A base of 64 snapshots is included with all MSA 1050/2050 systems without an additional license.) You can create snapshots immediately or schedule snapshot creation.

If the large pools feature is enabled, through use of the `large-pools` parameter of the `set advanced-settings` CLI command, the maximum number of volumes in a snapshot tree is limited to 9 (base volume plus 8 snapshots). The maximum number of volumes per snapshot will decrease to fewer than 9 if more than 3 replication sets are defined for volumes in the snapshot tree. If creating a snapshot will exceed the limit, you will be unable to create the snapshot unless you delete a snapshot first.

To see more information about a volume or snapshot, hover the cursor over the item in the volumes table.

You can view additional snapshot information by hovering the cursor over the snapshot in the Snapshots table. “[Viewing volumes](#)” (page 96) contains more details about the Volume Information and Snapshot Information panels that appear.

To create snapshots

1. In the Volumes topic, select from 1–16 volumes or snapshots.

NOTE: You can also select a combination of volumes and snapshots.

2. Select **Action > Create Snapshot**. The Create Snapshots panel opens.
3. Optional: In the Snapshot Name field, change the name for the snapshot. The default is *volume-name_sn*, where *n* starts at 0001. A snapshot name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system or include the following: " , < \
- If the name is used by another snapshot, you are prompted to enter a different name.
4. Optional: If you want to schedule a create-snapshot task, perform the following:
 - o Select the **Scheduled** check box.
 - o Optional: Change the default prefix to identify snapshots created by this task. The default is *volumesn*, where *n* starts at 01. The prefix is case sensitive and can have a maximum of 26 bytes. It cannot already exist in the system or include the following: " , < \
 - Scheduled snapshots are named *prefix_Sn*, where *n* starts at 0001.
 - o Optional: Select the number of snapshots to retain from either 1–8 if the large pools feature is enabled, or 1–16 if the large pools feature is disabled. The default is 1. When the task runs, the retention count is compared with the number of existing snapshots:
 - If the retention count has not been reached, the snapshot is created.
 - If the retention count has been reached, the oldest snapshot for the volume is unmapped, reset, and renamed to the next name in the sequence.
 - o Specify a date and a time at least five minutes in the future to run the task. The date must use the format *yyyy-mm-dd*. The time must use the format *hh:mm* followed by either **AM**, **PM**, or **24H** (24-hour clock). For example, 13:00 24H is the same as 1:00 PM.
 - o Optional: If you want the task to run more than once, perform the following:
 - Select the **Repeat** check box and specify how often the task should run.
 - Optional: Select the **End** check box to specify when the task should stop running.
 - Optional: Select the **Time Constraint** check box to specify a time range within which the task should run.
 - Optional: Select the **Date Constraint** check box to specify days when the task should run. Ensure that this constraint includes the start date.
5. Click **OK**.
 - o If **Scheduled** is not selected, the snapshot is created.
 - o If **Scheduled** is selected, the schedule is created and can be viewed in the Manage Schedules panel. For information on modifying or deleting schedules through this panel, see “[Managing scheduled tasks](#)” (page 59).

Resetting a snapshot

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 mappings are not changed.

This feature is supported for all snapshots in a tree hierarchy. However, a virtual snapshot can only be reset to the parent volume or snapshot from which it was created.

△ CAUTION: To avoid data corruption, unmount a snapshot from hosts before resetting the snapshot.

You can reset a snapshot immediately. You also have the option of scheduling a reset-snapshot task.

To see more information about a snapshot, hover the cursor over the item in the volumes table. You can view different snapshot information by hovering the cursor over the snapshot in the Snapshots table. [“Viewing volumes” \(page 96\)](#) contains more details about the Volume Information and Snapshot Information panels that appear.

To reset a snapshot

1. Unmount the snapshot from hosts.
2. In the Volumes topic, select a snapshot.
3. Select **Action > Reset Snapshot**. The Reset Snapshot panel opens.
4. Optional: To schedule a reset task, perform the following:
 - Select the **Schedule** check box.
 - Specify a date and a time at least five minutes in the future to run the task. The date must use the format `yyyy-mm-dd`. The time must use the format `hh:mm` followed by either **AM**, **PM**, or **24H** (24-hour clock). For example, 13:00 24H is the same as 1:00 PM.
 - Optional: If you want the task to run more than once:
 - Select the **Repeat** check box and specify how often the task should run.
 - Optional: Specify when the task should stop running.
 - Optional: Specify a time range within which the task should run.
 - Optional: Specify days when the task should run. Ensure that this constraint includes the start date.
5. Click **OK**. A confirmation panel appears.
6. Click **Yes** to continue. Otherwise, click **No**. If you clicked Yes:
 - If the **Schedule** check box was not selected, the snapshot is created. You can remount the snapshot.
 - If **Schedule** is selected, the schedule is created and can be viewed in the Manage Schedules panel, as described in [“Managing scheduled tasks” \(page 59\)](#). Make a reminder to unmount the snapshot before the scheduled task runs.

Creating a replication set from the Volumes topic

You can create a replication set, which specifies the components of a replication. The Create Replication Set panel enables you to create replication sets. You can access this panel from both the Replications and Volumes topics.

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

A peer connection must be defined to create and use a replication set. 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 volume group is part of a replication set, volumes cannot be added to or deleted from the volume group.

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

Primary volumes and volume groups

The volume, volume group, or snapshot that will be replicated is called the primary volume or volume group. It can belong to only one replication set. If the volume group is already in a replication set, individual volumes may not be included in separate replication sets. Conversely, if a volume that is a member of a volume group is already in a replication set, its volume group cannot be included in a separate replication set.

The maximum number of individual volumes and snapshots that can be replicated is 32 in total. If a volume group is being replicated, the maximum number of volumes that can exist in the group is 16.

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. 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.

For information on queuing replications between a system with MSA 1050/2050 controllers and a system with MSA 1040/2040 controllers, see [“Rules for using replication queue policy” \(page 34\)](#).

NOTE: 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.

Maintaining replication snapshot history from the Volumes topic

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

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

- You can set the number of snapshots to retain from 1–16, referred to as the snapshot retention count. This setting applies to management of snapshots for both the primary and secondary volume and can be changed at any time. Its value must be greater than the number of existing snapshots in the replication set, regardless of whether snapshot history is enabled. If you select a snapshot retention count value that is less than the current number of snapshots, an error message 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.
- The snapshots are named `basename_nnnn` where `_nnnn` starts at 0000 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.
- 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 existing 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.
- 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.
- A mapped snapshot history snapshot will not be deleted until after it is unmapped.
- 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.
- You can set the retention priority for snapshots to the following. In a snapshot tree, only leaf snapshots can be deleted automatically.
 - **never-delete.** Snapshots will never be deleted automatically to make space. The oldest snapshot in the snapshot history will be deleted once the snapshot count has been exceeded. This is the default.
 - **high.** Snapshots can be deleted after all eligible medium-priority snapshots have been deleted.
 - **medium.** Snapshots can be deleted after all eligible low-priority snapshots have been deleted.
 - **low.** Snapshots can be deleted. This parameter is unrelated to snapshot history, and because the default is never delete, snapshot history snapshots will normally not be affected in a low virtual memory situation.

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.

To create a replication set from the Volumes topic

1. In the volumes table, select a volume or snapshot to use as the primary volume.
2. Select **Action > Create Replication Set**. The Create Replication Set panel displays.
3. If the selected volume is in a volume group, source options appear.
 - To replicate the selected volume only, select **Single Volume**. This option is the default.
 - To replicate all volumes in the volume group, select **Volume Group**.
4. Enter a name for the replication set. The name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system, include leading or trailing spaces, or include the following characters: " , < \
5. Optional: Select a peer system to use as the secondary system for the replication set.
6. Optional: Select a pool on the secondary system. By default, the pool that corresponds with the pool in which the primary volume resides is selected. The selected pool must exist on the remote system.

7. Optional: If **Single Volume** is selected, enter a name for the secondary volume. The default name is the name of the primary volume. The name is case sensitive and can have a maximum of 32 bytes. It cannot already exist on the secondary system or include the following: " , < \
8. Optional: Specify the Queue Policy action to take when a replication is running and a new replication is requested.
9. Optional: Select the **Secondary Volume Snapshot History** check box to keep a snapshot history on the secondary system for the secondary volume.
 - o Set the **Retention Count** to specify the number of snapshots to retain.
 - o Modify the **Snapshot Basename** to change the snapshot name. The name is case sensitive and can have a maximum of 26 bytes. It cannot already exist in the system or include the following characters: " , < \
 - o Set the **Retention Priority** to specify the snapshot retention priority.
 - o Optional: Check **Primary Volume Snapshot History** to keep a snapshot history for the primary volume on the primary system
10. Optional: Select the **Scheduled** check box to schedule recurring replications.
11. Click **OK**.
12. In the success dialog box:
 - o If you selected the **Scheduled** check box, click **OK**. The Schedule Replications panel opens and you can set the options to create a schedule for replications. For more information on scheduling replications, see [“Initiating or scheduling a replication from the Volumes topic” \(page 108\)](#).
 - o Otherwise, you have the option to perform the first replication. Click **Yes** to begin the first replication, or click **No** to initiate the first replication later.

Initiating or scheduling a replication from the Volumes topic

After you have created a replication set, you can copy the selected volume or volume group on the primary system to the secondary system by initiating replication. The first time that you initiate replication, a full copy of the allocated pages for the volume or volume group 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 or create a scheduled task to initiate it automatically from both the Replications and Volumes topics. You can initiate replications only from a replication set's primary system. For information on modifying or deleting a replication schedule, see [“Managing replication schedules from the Volumes topic” \(page 110\)](#).

If a replication fails, 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.
- 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.

To manually initiate replication from the Volumes topic

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” \(page 124\)](#).

1. In the Volumes topic, select a replication set in the Replication Sets table.
2. Select **Action > Replicate**. The Replicate panel opens.
3. Click **OK**.
 - o If a replication is not in progress, the local system begins replicating the contents of the replication set volume to the remote system and the status of the replication set changes to **Running**.
 - o If a replication is already in progress, then the outcome of this replication request depends upon the Queue Policy setting specified in the Create Replication Set panel. For more information on setting the queue policy, see [“Queuing replications” \(page 106\)](#).

To schedule a replication from the Volumes topic

1. In the Volumes topic, select a replication set in the Replication Sets table.
2. Select **Action > Replicate**. The Replicate panel opens.
3. Select the **Schedule** check box.
4. Enter a name for the replication schedule task. The name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system or include the following: " , < \
5. Optional: If you want to create a replication of the last snapshot in the primary volume, select the **Last Snapshot** check box. At the time of the replication, the snapshot must exist. This snapshot may have been created either manually or by scheduling the snapshot. If no snapshot exists for the volume when the scheduled replication begins, the system generates an error with event code 362 and the replication fails.

NOTE: This option is unavailable when replicating volume groups.

6. Specify a date and a time in the future to be the first instance when the scheduled task will run, and to be the starting point for any specified recurrence.
 - o To set the **Date** value, enter the current date in the format **YYYY-MM-DD**.
 - o To set the **Time** value, enter two-digit values for the hour and minutes and select either **AM**, **PM**, or **24H** (24-hour clock). The minimum interval is one hour.
7. Optional: If you want the task to run more than once, select the **Repeat** check box.
 - o Specify how often the task should repeat. Enter a number and select the appropriate time unit. Replications can recur no less than 30 minutes apart.
 - o Either make sure the **End** check box is cleared, which allows the schedule to run indefinitely, or select the check box to specify when the schedule ends. To then specify an end date and time, select the **On** option, and specify when the schedule should stop running. Or, select the **After** option, and specify the number of replications that can occur before the schedule stops running.
 - o Either make sure the **Time Constraint** check box is cleared, which allows the schedule to run at any time, or select the check box to specify a time range within which the schedule should run.
 - o Either make sure the **Date Constraint** check box is cleared, which allows the schedule to run on any day, or select the check box to specify the days when the schedule should run.
8. Click **OK**. The schedule is created.

Managing replication schedules from the Volumes topic

You can modify or delete scheduled replication tasks on the primary system.

To manage a replication schedule from the Volumes topic

1. In the Volumes topic, select a replication set on the primary system that has an associated schedule from the Replication Sets table.
2. Select **Action > Manage Schedules**. The **Manage Schedules** panel opens.
3. Select the schedule to modify. Its settings display at the bottom of the panel.
4. If you want to create a replication of the last snapshot in the primary volume, select the **Last Snapshot** check box. At the time of the replication, the snapshot must exist. This snapshot may have been created either manually or by scheduling the snapshot. If no snapshot exists for the volume when the scheduled replication begins, event 362 will be logged and the replication fails

NOTE: This option is unavailable when replicating volume groups.

5. Specify a date and a time in the future to be the first instance when the scheduled task will run, and to be the starting point for any specified recurrence.
 - o To set the **Date** value, enter the current date in the format YYYY-MM-DD.
 - o To set the **Time** value, enter two-digit values for the hour and minutes and select either **AM**, **PM**, or **24H** (24-hour clock).
6. If you want the task to run more than once, select the **Repeat** check box.
 - o Specify how often the task should repeat. Enter a number and select the appropriate time unit. Replications can recur no less than 30 minutes apart.
 - o Either make sure the **End** check box is cleared, which allows the schedule to run without an end date, or select the check box and specify when the schedule should stop running.
 - o Either make sure the **Time Constraint** check box is cleared, which allows the schedule to run at any time, or select the check box to specify a time range within which the schedule should run.
 - o Either make sure the **Date Constraint** check box is cleared, which allows the schedule to run on any day, or select the check box to specify the days when the schedule should run.
7. Click **Apply**. A confirmation panel appears.
8. Click **Yes** to continue. Otherwise click **No**. If you clicked **Yes**, the schedule is modified.
9. Click **OK**.

To delete a schedule from the Volumes topic

1. In the Volumes topic, select a replication set on the primary system that has an associated schedule from the Replication Sets table.
2. Select **Action > Manage Schedules**. The **Manage Schedules** panel opens.
3. Select the schedule to delete.
4. Click **Delete Schedule**. A confirmation panel appears.
5. Click **Yes** to continue. Otherwise, click **No**. If you clicked **Yes**, the schedule was deleted.
6. Click **OK**.

7 Working in the Mappings topic

Viewing mappings

The Mapping topic shows a tabular view of information about mappings that are defined in the system. By default, the table shows 20 entries at a time and is sorted first by host and second by volume. For information about using tables, see [“Tips for using tables” \(page 12\)](#).

The mapping table shows the following information:

- Group.Host.Nickname. Identifies the initiators to which the mapping applies:
 - All Other Initiators. The mapping applies to all initiators that are not explicitly mapped with different settings.
 - *initiator-name*—The mapping applies to the initiator only.
 - *initiator-ID*—The mapping applies to the initiator only, and the initiator has no nickname.
 - *host-name.**—The mapping applies to all initiators in the host.
 - *host-group-name.**—The mapping applies to all hosts in this group.
- Volume. Identifies the volumes to which the mapping applies:
 - *volume-name*—The mapping applies to the volume only.
 - *volume-group-name.**—The mapping applies to all volumes in the volume group.
- Access. Shows the type of access assigned to the mapping:
 - *read-write*—The mapping permits read and write access to volumes.
 - *read-only*—The mapping permits read access to volumes.
 - *no-access*—The mapping prevents access to volumes.
- LUN. Shows whether the mapping uses a single LUN or a range of LUNs (indicated by *).
- Ports. Lists the controller host ports to which the mapping applies. Each number represents corresponding ports on both controllers.

To display more information about a mapping, see [“Viewing map details” \(page 114\)](#).

Mapping initiators and volumes

You can map initiators and volumes to control host access to volumes unless the volume is the secondary volume of a replication set. (Mapping also applies to hosts and host groups as well as initiators, and snapshots and volume groups as well as volumes. For the purposes of brevity, the terms *initiator* and *volumes* will stand in for all possibilities, unless otherwise stated.) By default, volumes are not mapped.

If a volume is mapped to All Other Initiators, this is its default mapping. The *default mapping* enables all connected initiators to see the volume using the specified access mode, LUN, and port settings. The advantage of a default mapping is that all connected initiators can discover the volume with no additional work by the administrator. The disadvantage is that all connected initiators can discover the volume with no restrictions. Therefore, this process is not recommended for specialized volumes that require restricted access. Also, to avoid multiple hosts mounting the volume and causing corruption, the hosts 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 corruption. To control access by specific hosts, you can create an *explicit mapping*. An explicit mapping can use different access mode, LUN, and port settings to allow or prevent access by a host to a volume, overriding the default mapping. When an explicit mapping is deleted, the volume's default mapping takes effect.

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 in 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 disk group ownership. When ULP is in use, 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.

If a group (host group or host) is mapped to a volume or volume group, all of the initiators within that group will have an individual map to each volume that makes up the request. As long as the group entity is mapped consistently, that set of individual maps will be represented as a grouped mapping. If any individual map within that group is modified, the grouped mapping will no longer be consistent, and it will no longer appear in the SMU. It will be replaced in the SMU with all of the individual maps.

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

NOTE: You cannot map LUN 0 for a SAS initiator. You can create a maximum of 1024 volumes, but because the supported LUN range is 1–1023, only 1023 volumes can be mapped using default mapping. Using explicit mapping, all volumes can be mapped.

NOTE: The secondary volume of a replication set cannot be mapped. Create a snapshot of the secondary volume or volume group and use the snapshot for mapping and accessing data.

To map initiators and volumes

1. Perform one of the following:
 - o In the Hosts topic, select the initiators to map and select **Action > Map Initiators**.
 - o In the Volumes topic, select the volumes to map and select **Action > Map Volumes**.
 - o In the Mapping topic, select **Map** to create a new mapping.
 - o In the Mapping topic, select one or more mappings to modify or delete and select **Action > Map**. You can also create a new mapping.

The Map panel opens and shows two tables side-by-side that list available initiators and volumes. You can use these tables to create mappings. There is also a table underneath the host and volume tables that lists mappings. After you create a mapping and before you save it, the mapping appears in the mappings table and you can modify its settings or delete it.

The Available Host Groups, Hosts, and Initiators table shows one or more of the following rows:

Table 12 Available host groups, hosts, and initiators

Row description	Group	Host	Nickname	ID
A row with these values always appears. Select this row to apply map settings to all initiators and create a default mapping.	-	-	(blank)	All Other Initiators
A row with these values appears for an initiator that is grouped into a host. Select this row to apply map settings to all initiators in this host.	-	<i>host-name</i>	*	*
A row with these values appears for an initiator that is grouped into a host group. Select this row to apply map settings to all initiators in this host group.	<i>host-group-name</i>	*	*	*
A row with these values appears for each initiator. Select this row to apply map settings to this initiator.	- or <i>host</i> - <i>host-group-name</i>	- or <i>host-name</i>	(blank) or <i>initiator-nick name</i>	<i>initiator-ID</i>

The Available Volume Groups and Volumes table shows one or more of the following rows:

Table 13 Available volume groups and volumes

Row description	Group	Name	Type
A row with these values appears for a volume/snapshot that is grouped into a volume group. Select this row to apply map settings to all volumes/snapshots in this volume group.	<i>volume-group-name</i>	*	Group
A row with these values appears for each volume/snapshot. Select this row to apply map settings to this volume/snapshot.	-	<i>volume-name</i>	<i>volume-type</i>

NOTE:

- When you select one or more host groups, hosts, or initiators in the Hosts topic, the item(s) appears in the Available Host Groups, Hosts, and Initiators table while all available volumes, volume groups, and snapshots appear in the Available Volume Groups and Volumes table.
- The converse is true when you select one or more volumes, volume groups, or snapshots in the Available Volume Groups and Volumes table.
- When you open the Map panel through the Mapping topic without selecting a mapping, both tables are fully populated with all available items.
- When you select a mapping in the mapping table, it appears in the list of mappings below the above two tables. Also, both tables are fully populated.

2. Perform one of the following:

- If nothing was pre-selected, select one or more initiators and one or more volumes to map and click the **Map** button.
- If initiators were pre-selected, select volumes to map to those initiators and click the **Map** button.
- If volumes were pre-selected, select initiators to map to those volumes and click the **Map** button.
- If maps were pre-selected, they already appear in the mapping table and a **Map** button will be displayed.


For each pairing of selected initiators and volumes, a row appears in the mapping table at the bottom of the panel. At this time, no further mappings can be added to the list. Mappings in the list can be modified — including the mapping's mode, LUN, or ports, or they can be deleted.

NOTE: Once a set of mappings between initiators and volumes have been defined using the **Map** button, the button changes from **Map** to **Reset**. If mappings have been pre-selected, the **Reset** button, not the **Map** button, appears.

3. Perform any of the following:

- To immediately remove a row from the table, in the Action column, select **Remove Row**.
- To delete an existing mapping, in the Action column, select **Delete**.
- To edit a mapping, set the following options:
 - **Mode.** The access mode can specify read-write access, read-only access, or no access to a volume. The default is read-write. When a mapping specifies no access, the volume is masked, which means it is not visible to associated initiators. Masking is useful to override an existing default map that allows open access so that access is denied only to specific initiators. To allow access to specific host(s) and deny access to all other hosts, create explicit map(s) to those hosts. For example, an engineering volume could be mapped with read-write access for the Engineering server and read-only access for servers used by other departments.

- **LUN.** The LUN identifies the volume to a host. The default is the lowest available LUN. Both controllers share one set of LUNs, and any unused LUN can be assigned to a mapping. However, each LUN is generally only used once as a default LUN. For example, if LUN 5 is the default for Volume1, no other volume in the storage system can use LUN 5 on the same port as its default LUN. For explicit mappings, the rules differ: LUNs used in default mappings can be reused in explicit mappings for other volumes and other hosts.

 **TIP:** When mapping a volume to a host with the Linux ext3 file system, specify read-write access. Otherwise, the file system will be unable to mount the volume and will report an error such as “unknown partition table.”

- **Ports.** Port selections specify controller host ports through which initiators are permitted to access, or are prevented from accessing, the volume. Selecting a port number automatically selects the corresponding port in each controller. By default, all ports are selected.
 - o To save a new mapping or edits to an existing mapping, in the Action column, select **Save**.
 - o To clear the mapping table and discard any changes, click **Reset**.
4. Once the list is correct, to apply changes, click **Apply** or **OK**. A confirmation panel appears. To discard the changes instead of applying them, click **Reset**.
 5. Click **Yes** to continue. Otherwise, click **No**. If you clicked Yes, the mapping changes are processed.
 6. To close the panel, click **Cancel**.

Removing mappings

You can remove one or more selected mappings between initiators and volumes.

To remove selected mappings from the system

1. Perform one of the following:
 - o In the Mapping topic, select one or more mappings from the table.
 - o In the Volumes topic, select at least one mapping in the Related Maps table.
2. Select **Action > Remove Mappings**. The Remove Mappings panel opens and the selected mappings display.
3. Click **OK**. The selected mappings are removed.

Removing all mappings

You can remove all mappings between initiators and volumes from the system.

To remove all mappings from the system

1. In the Mapping topic, select one or more mappings from the table.
2. Select **Action > Remove All Mappings**. The Remove All Mappings panel opens.
3. Click **OK**. The mappings are removed from the system.

Viewing map details

In the Hosts, Volumes, and Mapping topics, you can see basic information about mappings between hosts and volumes.

To view additional details

1. Perform one of the following:
 - o In the Hosts or Volumes topic, in the Related Maps table, select at least one mapping.
 - o In the Mapping topic, in the mapping table, select at least one mapping.

2. Select **Action > View Map Details**. The Map Details panel opens and shows the following information. For information about using tables, see [“Tips for using tables” \(page 12\)](#).
 - Host Group. Identifies the host group to which the mapping applies:
 - -. The mapping does not apply to a host group.
 - *host-group-name*. The mapping applies to all hosts in this host group.
 - Host. Identifies the host to which the mapping applies:
 - -. The mapping does not apply to a host.
 - *host-name*. The mapping applies to all initiators in this host.
 - Nickname. Shows the nickname of the initiator, if a nickname is assigned. Otherwise, this field is blank.
 - Initiator ID. Shows the WWN of an FC or SAS initiator or the IQN of an iSCSI initiator.
 - Volume Group. Identifies the volumes to which the mapping applies:
 - -. The mapping does not apply to a volume group.
 - *volume-group-name*. The mapping applies to all volumes in this volume group.
 - Volume. Identifies the volume to which the mapping applies.
 - Access. Shows the type of access assigned to the mapping:
 - *read-write*—The mapping permits read and write access to volumes.
 - *read-only*—The mapping permits read access to volumes.
 - *no-access*—The mapping prevents access to volumes.
 - LUN. Shows whether the mapping uses a single LUN or a range of LUNs (indicated by *). By default, the table is sorted by this column.
 - Ports. Lists the controller host ports to which the mapping applies. Each number represents corresponding ports on both controllers.
3. Click **OK**.

8 Working in the Replications topic

About replicating virtual volumes

Replication for virtual storage is a licensed feature that provides a remote copy of a volume, volume group, or snapshot (hereafter known as *volume*) 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 using all pages of data from the source.) Remote Snap can be used for disaster recovery, to preserve data, and to back data up to off-site locations. It can also be used to distribute data.

Replication prerequisites

To replicate a volume, you must first create a peer connection and 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 a replication license for virtual storage. 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, use the `query port-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 documentation. For more information on peer connections, see [“Creating a peer connection” \(page 123\)](#), [“Deleting a peer connection” \(page 126\)](#), and [“Modifying a peer connection” \(page 125\)](#).

After you create a peer connection, you can create a replication set. A replication set specifies a volume, snapshot, or multiple volumes in a volume group (hereafter known as *volume*) 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, a corresponding volume is automatically created on the other system of the peer connection, known as the secondary system, along with the infrastructure needed for replication. The infrastructure consists of internal snapshots used for replication operations:

- A replication set for a volume consumes two internal snapshots each for the primary volume and the secondary volume if the queue policy is set to `Discard`, or three each if the queue policy is set to `Queue Latest`.
- A replication set for a volume group consumes two internal volume groups if the queue policy is set to `Discard`, or three if the queue policy is set to `Queue Latest`. Each internal volume group contains a number of volumes equal to the number of volumes in the base volume group.

Using a volume group for a replication set enables you to make sure that multiple volumes are synchronized at the same time. When a volume group is replicated, snapshots of all of the volumes are created simultaneously. In doing so, it 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, 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 [“Creating a replication set from the Replications topic” \(page 126\)](#), [“Creating a replication set from the Volumes topic” \(page 105\)](#), [“Modifying a replication set” \(page 129\)](#), and [“Deleting a replication set” \(page 130\)](#).

NOTE: HPE recommends 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.

Subsequent replications are completed by resetting one of the hidden snapshots to contain the contents last replicated and then resetting the other hidden snapshot to the current primary volume contents and comparing the changes. The system writes any changes it finds on the hidden primary snapshot to the hidden secondary snapshot, after which the secondary volume is updated to contain the contents of the secondary volume.

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” \(page 132\)](#), [“Initiating or scheduling a replication from the Replications topic” \(page 130\)](#), [“Initiating or scheduling a replication from the Volumes topic” \(page 108\)](#), [“Resuming a replication” \(page 133\)](#), and [“Suspending a replication” \(page 132\)](#).

You can initiate a replication manually or by using a schedule. When creating a schedule for a replication set, you cannot specify for replication to occur more frequently than once an hour. For more information on scheduling a replication set, see [“Initiating or scheduling a replication from the Replications topic” \(page 130\)](#) and [“Initiating or scheduling a replication from the Volumes topic” \(page 108\)](#).

Initial replication

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

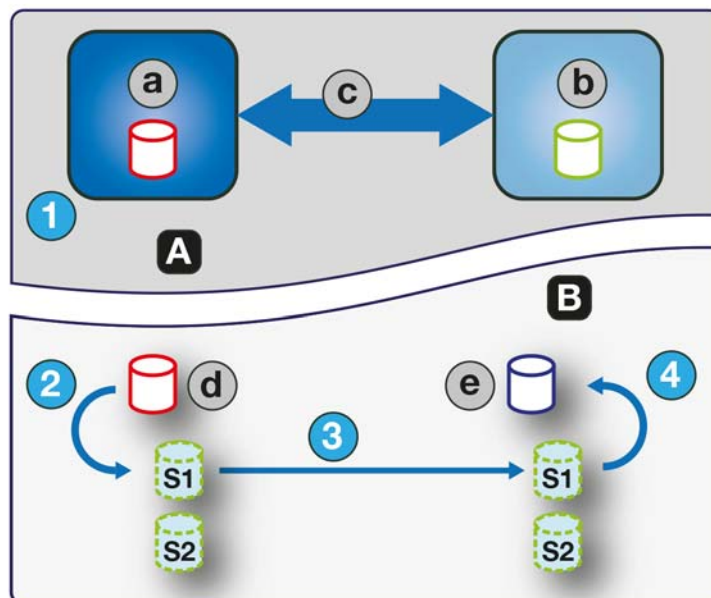


Figure 2 Replication process for initial replication

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 initiates replication on the primary system (step 1). The current primary volume contents, which might be different than when the replication set was created, replace the contents of S1 on the primary system (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 take place in replications subsequent to the initial replication of a single volume.

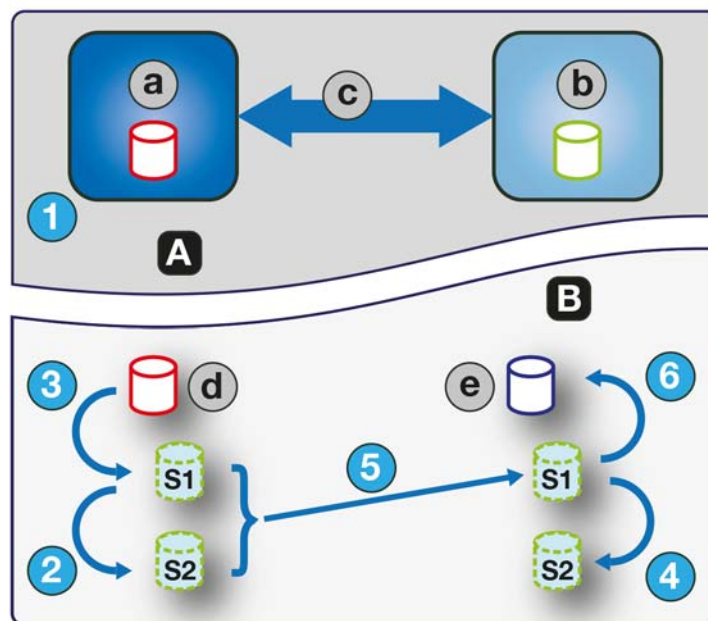


Figure 3 Replication process for replications subsequent to the initial replication.

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 initiates replication on the primary system (step 1). On the primary system, the S1 contents replace the S2 contents (step 2). (The S2 contents can then be used for comparison during step 5.) The current primary volume contents replace the contents of S1 on the primary system (step 3). On the secondary system, the S1 contents replace the S2 contents (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 three 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). 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

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. The policy will either notify you via the event log that the percentage has been reached (in which case the system continues to take snapshots, using the general pool space), or 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 on setting snapshot retention priority, see [“Maintaining replication snapshot history from the Replications topic” \(page 127\)](#).

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

If you are using Remote Snap and snapshot space management, there are specific 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 snapshot space to accommodate the maximum anticipated size of the two internal snapshots, which cannot be deleted, and any other snapshots that you would like to retain.
- 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 documentation.
- You can later create more snapshot space by adding disks to the pool to increase its size.

If the internal snapshots are larger than anticipated and take up a lot of snapshot space, you can adjust the snapshot space thresholds or increase the snapshot space to prevent unintentional automatic deletion of snapshots that you want to retain. 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 documentation.

Replication and empty allocated pages

Deleting data from a volume can result in deallocation of pages on that volume. 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

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

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

Accessing the data while keeping the replication set intact

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

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

To temporarily access data at the backup site

1. Create a snapshot of the secondary volume or use a snapshot history snapshot.
2. Map the snapshot to hosts.
3. When the data center system has recovered, delete the snapshot.

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

If you do not think the data center system can be recovered in time or at all, then you will want to temporarily access the data from the backup system as if it were the primary system. You can again create a snapshot of the secondary volume and map that to hosts, or delete the replication set to allow mapping the secondary volume directly to hosts. Deleting the replication set means the secondary volume becomes a base volume and is no longer the target of a replication. Should the primary volume become available and you want to use it as is in preparation for another disaster, a new replication set with a new secondary volume must be created. Deleting the replication set also enables cleaning up any leftover artifacts of the replication set.

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

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

Disaster recovery procedures

In a disaster recovery situation, you might typically:

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

To manually transfer operations from the data center system to the backup system

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

To restore operations to the data center system

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

To prepare the backup system for disaster recovery after the replication is complete

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

Replication licensing




For information about viewing the status of licensed features in your system, see [“Viewing the status of licensed features” \(page 49\)](#).

Viewing replications

The Replications topic shows a tabular view of information about peer connections, replication sets, and snapshot history of local snapshots associated with a selected replication set. For information about using tables, see [“Tips for using tables” \(page 12\)](#). For more information about replication, see [“About replicating virtual volumes” \(page 34\)](#).

Peer Connections table

The Peer Connections table shows the following information. By default, the table shows 10 entries at a time.

- Name. Shows the specified peer connection name.
- Status. Shows the status of the peer connection:
 - Online—The systems have a valid connection.
 - Offline—No connection is available to the remote system.
- Health. Shows the health of the component:  OK,  Fault, or  Unknown.
- Type. Shows the type of host ports being used for the peer connection: FC or iSCSI.
- Local Ports. Shows the IDs of host ports in the local system.
- Remote Ports. Shows the IDs of host ports in the remote system.



To see more information about a peer connection, hover the cursor over the peer connection in the table. The **Peer Connections** panel that appears contains the following information. If the health is not OK, the health reason and recommended action are shown to help you resolve problems.

Peer Connections	Name, serial number, connection type, connection status, local host port name and IP address, remote host port name and IP address, health
------------------	--

Replication Sets table

The Replication Sets table shows the following information. By default, the table shows 10 entries at a time.

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 system to enable management interfaces to show proper time values for replication operations.

- Name. Shows the replication set name.
- Primary Volume. Shows the primary volume name. For replication sets that use volume groups, the primary volume name is `volume-group-name.*` where `.*` signifies that the replication set contains more than one volume. If the volume is on the local system, the  icon appears.
- Secondary Volume. Shows the secondary volume name. For replication sets that use volume groups, the secondary volume name is `volume-group-name.*` where `.*` signifies that the replication set contains more than one volume. If the volume is on the local system, the  icon appears.
- Status. Shows the status of the replication set.
 - Not Ready—The replication set is not ready for replications because the system is still preparing the replication set.
 - Unsynchronized—The primary and secondary volumes are unsynchronized because the system has prepared the replication set, but the initial replication has not run.
 - Running—A replication is in progress.
 - Ready—The replication set is ready for a replication.
 - Suspended—Replications have been suspended.
 - Unknown—This system cannot communicate with the primary system and thus cannot be sure of the current state of the replication set. Check the state of the primary system.
- Last Successful Run. Shows the date and time of the last successful replication.
- Estimated Completion Time. Shows the estimated date and time for the replication in progress to complete.

To see more information about a replication set, hover the cursor over a replication set in the Replication Sets table. The **Replication Sets** panel that appears contains the following information.

Replication Sets	Replication set name and serial number; status; primary volume or volume group name and serial number; secondary volume or volume group name and serial number; peer connection name; queue policy, queue count, secondary volume snapshot history, primary volume snapshot history, retention count, retention priority, snapshot basename, associated schedule name, current run progress, current run start time, current run estimated time to completion, current run transferred data, last successful run, last run start time, last run end time, last run transferred data, last run status, and last run error status
------------------	---

Replication Snapshot History table

The Replication Snapshot History table shows the following information. By default, the table shows 10 entries at a time.

- Local Snapshot Name. Shows the local snapshot name.
- Creation Date/Time. Shows the date and time of the last successful snapshot created.

- Snap Data. Shows the total amount of write data associated with the snapshot.
- Unique Data. Shows the amount of write data that is unique to the snapshot.

To see more information about a snapshot history, hover the cursor over a snapshot set in the Replication Snapshot History table. The Snapshot Information hover panel that appears contains the following information.

Replication Snapshot History	Name, serial number, status, status reason, retention priority, snapshot data, unique data, shared data, pool, class, number of snaps, number of snapshots in tree, source volume, total size, creation date/time, type, parent volume, base volume, health
------------------------------	---

Querying a peer connection

You can view information about systems you might use in a peer connection before creating the peer connection, or you can view information about systems currently in a peer connection before modifying the peer connection.

To query a peer connection

1. In the Replications topic, do one of the following to display the Query Peer Connection panel:
 - o Select the peer connection to query in the Peer Connections table, then select **Action > Query Peer Connection**. The remote host port address field is pre-populated with the selected peer's remote port address.
 - o Select **Action > Query Peer Connection**.
2. If you did not select a peer connection from the Peer Connections table, enter the remote host port address to query in the text box.
3. Click **OK**. A processing dialog box appears while the remote port address is queried. If successful, detailed information about the remote system and controllers displays. An error message displays if the operation is unsuccessful.

Creating a peer connection

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 are not used for 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 ports. Ports at both ends of the connection must use the same protocol.
- Each system must have a virtual pool.
- If iSCSI CHAP is configured for the peer connection, the authentication must be valid.
- You must specify the username and password of a user with the manage role on the remote system.

NOTE: For information on creating a peer connection between a system with MSA 1050/2050 controllers and a system with MSA 1040/2040 controllers, see [“Replicating between MSA 1050/2050 and MSA 1040/2040 systems” \(page 34\)](#).

You can create a maximum of one peer connection per storage system for the MSA 1050 system, and four peer connections per storage system for the MSA 2050 system. However, only one peer connection is allowed to a particular remote system. Attempting to create a second peer connection to the same system will fail.

While creating the peer connection, the local system receives information about all host ports and IPs on the remote system as well as the remote system's licensing and host port health. (Both systems must be licensed to use Remote

Snap for virtual storage.) It also links host ports of the select host port type on the local system to those 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 the bi-directional communication path between the systems when exchanging information and transferring replicated data. Once you create a peer connection, you can use it when creating any replication set. Because the peer connection is bi-directional, replication sets can be created from both systems with replication occurring from either direction.

NOTE: You can use the `query peer-connection` CLI command to determine if the remote system is compatible with your system. This command provides information about the remote system, such as host ports, licensing, and pools. You can run it before creating the peer connection to determine if either system needs to be reconfigured first. You can also run it to diagnose problems if creating a peer connection fails.

To create a peer connection

1. In the Replications topic, select **Action > Create Peer Connection**. The Create Peer Connection panel opens.
2. Enter a name for the peer connection. The name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system or include the following: " , < \
3. Enter the destination port address for the remote system.
4. Enter the name and password of a user with the manage role on the remote system.
5. Click **OK**.
6. If the task succeeds, click **OK** in the confirmation dialog. The peer connection is created and the Peer Connections table is updated.

If the task does not succeed, the Create Peer Connection panel displays with errors in red text. Correct the errors, then click **OK**. For more information about system errors, see the Event Descriptions Reference Guide.

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. Peer connections support one-way CHAP only.

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, 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 (see the following procedure).

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

If you want to use Challenge Handshake Authentication Protocol (CHAP) for the iSCSI connection between peer systems, see the following procedure to set up CHAP. In a peer connection, both systems will alternately act as an 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.
 - o If both succeed, you can create, set, or perform replication on that peer connection.
 - o 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.

Modifying a peer connection

You can 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 configurations. 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. For information on modifying a peer connection between a system with MSA 1050/2050 controllers and a system with MSA 1040/2040 controllers, see [“Replicating between MSA 1050/2050 and MSA 1040/2040 systems” \(page 34\)](#).

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” \(page 124\)](#).

To modify a peer connection

1. In the Replications topic, select the peer connection to be modified in the Peer Connections table.
2. Select **Action > Modify Peer Connection**. The Modify Peer Connection panel displays.
3. Change one of the following (you cannot change both):
 - o Select **New Name**, then enter a new name for the peer connection. The name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system or include the following: " , < \
 - o Select **New Remote Address (FC-WWN or iSCSI-IP)**, then enter a new address for the remote system.

NOTE: You can change protocols used in the peer connection between FC and iSCSI for an MSA 2050 system by modifying the peer connection to use the remote port address of the new protocol.

4. Enter the name and password of a user assigned a manage role on the remote system.
5. Click **OK**. The peer connection is modified and the Peer Connections table is updated.

Deleting a peer connection

You can delete a peer connection if there are no replication sets that belong to the peer connection. If there are replication sets that belong to the peer connection, you must delete them before you can delete the peer connection. For more information, see [“Deleting a replication set” \(page 130\)](#).

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” \(page 124\)](#).

To delete a peer connection

1. In the Replications topic, select the peer connection to be deleted in the Peer Connections table.
2. Select **Action > Delete Peer Connection**.
3. Click **OK**. The peer connection is deleted and the Peer Connections table is updated.

Creating a replication set from the Replications topic

You can create a replication set, which specifies the components of a replication. The Create Replication Set panel enables you to create replication sets. You can access this panel from both the Replications and Volumes topics.

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

A peer connection must be defined to create and use a replication set. 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 volume group is part of a replication set, volumes cannot be added to or deleted from the volume group.

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

Primary volumes and volume groups

The volume, volume group, or snapshot that will be replicated is called the primary volume or volume group. It can belong to only one replication set. If the volume group is already in a replication set, individual volumes may not be included in separate replication sets. Conversely, if a volume that is a member of a volume group is already in a replication set, its volume group cannot be included in a separate replication set.

The maximum number of individual volumes and snapshots that can be replicated is 32 in total. If a volume group is being replicated, the maximum number of volumes that can exist in the group is 16.

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. 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.

For information on queuing replications between a system with MSA 1050/2050 controllers and a system with MSA 1040/2040 controllers, see [“Rules for using replication queue policy” \(page 34\)](#).

NOTE: 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.

Maintaining replication snapshot history from the Replications topic

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

- For a primary volume, when a replication starts it will create a snapshot of the data image being replicated.
- For a secondary volume, when a replication successfully completes it will create a snapshot of the data image just transferred to the secondary volume. (This is in contrast to the primary volume snapshot, which is created before the sync.) If replication does not complete, a snapshot will not be created.
- You can set the number of snapshots to retain from 1–16, referred to as the snapshot retention count. This setting applies to management of snapshots for both the primary and secondary volume and can be changed at any time. Its value must be greater than the number of existing snapshots in the replication set, regardless of whether snapshot history is enabled. If you select a snapshot retention count value that is less than the current number of snapshots, an error message 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.
- The snapshots are named `basename_nnnn` where `_nnnn` starts at 0000 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.
- 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 existing 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.
- You can set the retention priority for snapshots to the following. In a snapshot tree, only leaf snapshots can be deleted automatically.
 - **never-delete.** Snapshots will never be deleted automatically to make space. The oldest snapshot in the snapshot history will be deleted once the snapshot count has been exceeded. This is the default.
 - **high.** Snapshots can be deleted after all eligible medium-priority snapshots have been deleted.
 - **medium.** Snapshots can be deleted after all eligible low-priority snapshots have been deleted.
 - **low.** Snapshots can be deleted. This parameter is unrelated to snapshot history, and because the default is never delete, snapshot history snapshots will normally not be affected in a low virtual memory situation.

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.

To create a replication set from the Replications topic

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” \(page 124\)](#).

1. In the Peer Connections table, select the peer connection to use for the replication set.
2. Select **Action > Create Replication Set**. The Create Replication Set panel displays.
3. Enter a name for the replication set. The name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system, include leading or trailing spaces, or include the following characters: " , < \
4. Select whether you want to use a single volume or a volume group, which will filter the entries in the adjacent table.
5. In the table, select the volume or volume group to replicate. This will be the primary volume or volume group.
6. Optional: If **Single Volume** is selected, enter a name for the secondary volume. The default name is the name of the primary volume. The name is case sensitive and can have a maximum of 32 bytes. It cannot already exist on the secondary system or include the following: " , < \
7. Optional: Select a pool on the secondary system. By default, the pool that corresponds with the pool in which the primary volume resides is selected. The selected pool must exist on the remote system.
8. Optional: Specify the Queue Policy action to take when a replication is running and a new replication is requested.
9. Optional: Select the **Secondary Volume Snapshot History** check box to keep a snapshot history on the secondary system for the secondary volume.
 - Set the Retention Count to specify the number of snapshots to retain.
 - Modify the Snapshot Basename to change the snapshot name. The name is case sensitive and can have a maximum of 26 bytes. It cannot already exist in the system or include the following characters: " , < \
 - Set the Retention Priority to specify the snapshot retention priority.
 - Optional: Check **Primary Volume Snapshot History** to keep a snapshot history for the primary volume on the primary system.
10. Optional: Select the **Scheduled** check box to schedule recurring replications.

11. Click **OK**.
12. In the success dialog box:
 - o If you selected the **Scheduled** check box, click **OK**. The Schedule Replications panel opens and you can set the options to create a schedule for replications. For more information on scheduling replications, see [“Initiating or scheduling a replication from the Replications topic” \(page 130\)](#).
 - o Otherwise, you have the option to perform the first replication. Click **Yes** to begin the first replication, or click **No** to initiate the first replication later.

Modifying a replication set

You can change a replication set's name, queue policy, and snapshot history settings. Volume membership of a replication cannot change for the life of 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” \(page 124\)](#).

To modify a replication set

1. In the Replications topic, select the replication set in the Replications Sets table that you want to modify.
2. Select **Action > Modify Replication Set**. The Modify Replication Set panel opens.
3. Enter a new name for the replication set. The name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system, include leading or trailing spaces, or include the following: " , < \
4. Specify the Queue Policy action to take when a replication is running and a new replication is requested.
 - o **Discard**. Discard the new replication request.
 - o **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. 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.
5. Optional: Select the **Secondary Volume Snapshot History** check box to keep a snapshot history on the secondary system for the secondary volume.
 - o Set the Retention Count to modify the number of snapshots to retain. Its value must be greater than the number of existing snapshots in the replication set, regardless of whether snapshot history is enabled.

NOTE: 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.

- o Set the Snapshot Basename to modify the snapshot name. The name is case sensitive and can have a maximum of 26 bytes. It cannot already exist in the system or include the following characters: " , < \

NOTE: If you change the Snapshot Basename 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.

- o Set the Retention Priority to specify the snapshot retention priority.
 - o Optional: Check **Primary Volume Snapshot History** to keep a snapshot history for the primary volume on the primary system.
6. Click **OK**. The name of the replication set is updated in the Replications Sets table.

Deleting a replication set

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

If you want to delete a replication set that has a replication in progress, you must first suspend and then abort replication for that replication set. For more information, see [“Aborting a replication” \(page 132\)](#) or [“Suspending a replication” \(page 132\)](#). When a replication set is deleted, the snapshots created from the snapshot history feature will not be deleted. You will be able to manage those snapshots like any other snapshots. For more information, see [“Maintaining replication snapshot history from the Replications topic” \(page 127\)](#).

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.

To delete a replication set

1. In the Replications topic, select the replication set to be deleted in the Replication Sets table.
2. Select **Action > Delete Replication Set**.
3. Click **OK**. The replication set is deleted and the Replication Sets table is updated.

Initiating or scheduling a replication from the Replications topic

After you have created a replication set, you can copy the selected volume or volume group on the primary system to the secondary system by initiating replication. The first time that you initiate replication, a full copy of the allocated pages for the volume or volume group 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 or create a scheduled task to initiate it automatically from both the Replications and Volumes topics. You can initiate replications only from a replication set's primary system.

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 system to enable management interfaces to show proper time values for replication operations.

If a replication fails, 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.
 - 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.

To manually initiate replication from the Replications topic

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” \(page 124\)](#).

1. In the Replications topic, select a replication set in the Replication Sets table.
2. Select **Action > Replicate**. The Replicate panel opens.
3. Click **OK**.
 - If a replication is not in progress, the local system begins replicating the contents of the replication set volume to the remote system and the status of the replication set changes to *Running*.
 - If a replication is already in progress, then the outcome of this replication request depends upon the Queue Policy setting specified in the Create Replication Set panel. For more information on setting the queue policy, see [“Queuing replications” \(page 127\)](#).

To schedule a replication from the Replications topic

1. In the Replications topic, select a replication set from the Replication Sets table.
2. Select **Action > Replicate**. The Replicate panel opens.
3. Select the **Schedule** check box.
4. Enter a name for the replication schedule task. The name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system or include the following: " , < \
5. Optional: If you want to create a replication of the last snapshot of the primary volume, select the **Last Snapshot** check box. At the time of the replication, the snapshot must exist. This snapshot may have been created either manually or by scheduling the snapshot. If no snapshot exists for the volume when the scheduled replication begins, event 362 will be logged and the replication fails.

NOTE: This option is unavailable when replicating volume groups.

6. Specify a date and a time in the future to be the first instance when the scheduled task will run, and to be the starting point for any specified recurrence.
 - To set the **Date** value, enter the current date in the format *YYYY-MM-DD*.
 - To set the **Time** value, enter two-digit values for the hour and minutes and select either **AM**, **PM**, or **24H** (24-hour clock). The minimum interval is one hour.

7. If you want the task to run more than once, select the **Repeat** check box.
 - o Specify how often the task should repeat. Enter a number and select the appropriate time unit. Replications can recur no less than 30 minutes apart.
 - o Either make sure the **End** check box is cleared, which allows the schedule to run indefinitely, or select the check box to specify when the schedule ends. To then specify an end date and time, select the **On** option, and specify when the schedule should stop running. Or, select the **After** option, and specify the number of replications that can occur before the schedule stops running.
 - o Either make sure the **Time Constraint** check box is cleared, which allows the schedule to run at any time, or select the check box to specify a time range within which the schedule should run.
 - o Either make sure the **Date Constraint** check box is cleared, which allows the schedule to run on any day, or select the check box to specify the days when the schedule should run.
8. Click **OK**. The schedule is created.

Aborting a replication

You can abort running or suspended replication operations for a specified replication set, only from its primary system. Aborting a replication for a replication set that is in a **Ready** or **Unsynchronized** state will generate an error.

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.

To abort a replication

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” \(page 124\)](#).

1. In the Replications topic, select a replication set that is currently being replicated in the Replication Sets table.
2. Select **Action > Abort Replication**.
3. Click **OK**. The replication is aborted.

Suspending a replication

You can suspend replication operations for a specified replication set from its primary system. You can suspend replications from a replication set's primary system only.

When you suspend a replication set, all replications in progress are paused and no new replications are allowed to occur. You can abort suspended replications. After you suspend replication, you must resume it to allow the replication set to resume replications that were in progress and allow new replications to occur. For more information, see [“Aborting a replication” \(page 132\)](#) or [“Resuming a replication” \(page 133\)](#).

If replications are attempted during the suspended period (including scheduled replications), the replications will fail.

To suspend a replication

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” \(page 124\)](#).

1. In the Replications topic, select a replication set that is currently being replicated in the Replication Sets table.
2. Select **Action > Suspend Replication**.
3. Click **OK**. The replications on the replication set are suspended and the status of the replication set changes to *Suspended*.

Resuming a replication

You can resume the replication operations of a specified suspended replication set. You can resume replications from a replication set's primary system only.

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. If you aborted a replication while the replication set was suspended, the aborted replication does not resume.

To resume a replication

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” \(page 124\)](#).

1. In the Replications topic, select a replication set for which replications were suspended in the Replication Sets table.
2. Select **Action > Resume Replication**.
3. Click **OK**. Replications on the replication set are resumed and the status of the replication set changes to *Running*.

Managing replication schedules from the Replications topic

You can modify or delete scheduled replication tasks on the primary system.

To manage a replication schedule

1. In the Replications topic, select a replication set on the primary system that has an associated schedule from the Replication Sets table.
2. Select **Action > Manage Schedules**. The **Manage Schedules** panel opens.
3. Select the schedule to modify. Its settings display at the bottom of the panel.
4. If you want to create a replication of the last snapshot in the primary volume, select the **Last Snapshot** check box.

NOTE: This option is unavailable when replicating volume groups.

5. Specify a date and a time in the future to be the first instance when the scheduled task will run, and to be the starting point for any specified recurrence.
 - o To set the **Date** value, enter the current date in the format YYYY-MM-DD.
 - o To set the **Time** value, enter two-digit values for the hour and minutes and select either **AM**, **PM**, or **24H** (24-hour clock).

6. If you want the task to run more than once, select the **Repeat** check box.
 - o Specify how often the task should repeat. Enter a number and select the appropriate time unit. Replications can recur no less than 30 minutes apart.
 - o Either make sure the **End** check box is cleared, which allows the schedule to run without an end date, or select the check box and specify when the schedule should stop running.
 - o Either make sure the **Time Constraint** check box is cleared, which allows the schedule to run at any time, or select the check box to specify a time range within which the schedule should run.
 - o Either make sure the **Date Constraint** check box is cleared, which allows the schedule to run on any day, or select the check box to specify the days when the schedule should run.
7. Click **Apply**. A confirmation panel appears.
8. Click **Yes** to continue. Otherwise click **No**. If you clicked Yes, the schedule is modified.
9. Click **OK**.

To delete a schedule from the Replications topic

1. In the Replications topic, select a replication set on the primary system that has an associated schedule from the Replication Sets table.
2. Select **Action > Manage Schedules**. The **Manage Schedules** panel opens.
3. Select the schedule to delete.
4. Click **Delete Schedule**. A confirmation panel appears.
5. Click **Yes** to continue. Otherwise, click **No**. If you clicked Yes, the schedule was deleted.
6. Click **OK**.


9 Working in the Performance topic

Viewing performance statistics

The Performance topic shows performance statistics for the following types of components: disks, disk groups, virtual pools, virtual tiers, host ports, controllers, and volumes. For more information about performance statistics, see [“About performance statistics” \(page 29\)](#).

You can view current statistics in tabular format for all component types, and historical statistics in graphical format for disks, disk groups, and virtual pools and tiers.

To view performance statistics

1. In the Performance topic, select a component type from the Show list. The components table shows information about each component of that type in the system. For information about using tables, see [“Tips for using tables” \(page 12\)](#).
 2. Select one or more components in the list.
 3. Click **Show Data**. The Current Data area shows the sample time, which is the date and time when the data sample was collected. It also shows the total duration of all data samples, which is the time period between collection and display of the current sample, the previous sample (if any), and a table of current performance statistics for each selected component.
 4. To view graphs of historical data for the selected disks, disk groups, virtual pools, or virtual tiers, select the **Historical Data** check box. The Historical Data area shows the time range of samples whose data is represented by the graphs, and the Total IOPS graph by default.
 5. To specify either a time range or a count of historical statistics samples to display, perform the following:
 - o Click **Set time range**. The Update Historical Statistics panel opens and shows the default count value of 100.
 - o To specify a count, in the Count field, enter a value in the range of 5–100 and click **OK**.
 - o To specify a time range, perform the following:
 - Select the **Time Range** check box.
 - Set date/time values for the starting and ending samples. The values must be between the current date/time and 6 months in the past. The ending values must be more recent than the starting values.
-
-  **TIP:** If you specify a time range, it is recommended to specify a range of 24 hours or less.
-
- Click **OK**. In the Historical Data area, the Time Range values are updated to show the times of the oldest and newest samples displayed, and the graph for the selected components is updated.
6. To view different historical statistics, select a graph from the Statistics list. For a description of each graph, see [“Historical performance graphs” \(page 135\)](#).
 7. To hide the legend in the upper right corner of a historical statistics graph, clear the **Show Legend** check box.

Historical performance graphs

The following table describes the graphs of historical statistics that are available for each component type. In the graphs, measurement units are automatically scaled to best represent the sample data within the page space.

Table 14 Historical performance graphs

System component	Graph	Description
Disk, group, pool, tier	Total IOPS	Shows the total number of read and write operations per second since the last sampling time.
Disk, group, pool, tier	Read IOPS	Shows the number of read operations per second since the last sampling time.

Table 14 Historical performance graphs (continued)

System component	Graph	Description
Disk, group, pool, tier	Write IOPS	Shows the number of write operations per second since the last sampling time.
Disk, group, pool, tier	Data Throughput	Shows the overall rate at which data was read and written since the last sampling time.
Disk, group, pool, tier	Read Throughput	Shows the rate at which data was read since the last sampling time.
Disk, group, pool, tier	Write Throughput	Shows the rate at which data was written since the last sampling time.
Disk, group, pool, tier	Total I/Os	Shows the number of read and write operations since the last sampling time.
Disk, group, pool, tier	Number of Reads	Shows the number of read operations since the last sampling time.
Disk, group, pool, tier	Number of Writes	Shows the number of write operations since the last sampling time.
Disk, group, pool, tier	Data Transferred	Shows the total amount of data read and written since the last sampling time.
Disk, group, pool, tier	Data Read	Shows the amount of data read since the last sampling time.
Disk, group, pool, tier	Data Written	Shows the amount of data written since the last sampling time.
Disk, group	Average Response Time	Shows the average response time for reads and writes since the last sampling time.
Disk, group	Average Read Response Time	Shows the average response time for reads since the last sampling time.
Disk, group	Average Write Response Time	Shows the average response time for writes since the last sampling time.
Disk, group	Average I/O Size	Shows the average size of reads and writes since the last sampling time.
Disk, group	Average Read I/O Size	Shows the average size of reads since the last sampling time.
Disk, group	Average Write I/O Size	Shows the average size of writes since the last sampling time.
Disk, group	Number of Disk Errors	Shows the number of disk errors since the last sampling time.
Disk, group	Queue Depth	Shows the average number of pending I/O operations being serviced since the last sampling time. This value represents periods of activity only and excludes periods of inactivity.
Pool, tier	Number of Allocated Pages	Shows the number of 4-MB pages allocated to volumes, based on writes to those volumes. Creating a volume does not cause any allocations. Pages are allocated as data is written.
Tier	Number of Page Moves In	Shows the number of pages moved into this tier from a different tier.
Tier	Number of Page Moves Out	Shows the number of pages moved out of this tier to other tiers.
Tier	Number of Page Rebalances	Shows the number of pages moved between disk groups in this tier to automatically load balance.

Table 14 Historical performance graphs (continued)

System component	Graph	Description
Tier	Number of Initial Allocations	Shows the number of pages that are allocated as a result of host writes. This number does not include pages allocated as a result of background tiering page movement. (Tiering moves pages from one tier to another, so one tier will see a page deallocated, while another tier will show pages allocated; these background moves are not considered “initial allocations.”)
Tier	Number of Unmaps	Shows the number of 4-MB pages that are automatically reclaimed and deallocated because they are empty (they contain only zeroes for data).
Tier	Number of RFC Copies	Shows the number of 4-MB pages copied from spinning disks to SSD read cache (read flash cache).
Tier	Number of Zero-Pages Reclaimed	Shows the number of empty (zero-filled) pages that were reclaimed during this sample period.

Updating historical statistics

The Performance topic can show historical performance statistics for the following types of components: disks, disk groups, and virtual pools and tiers. By default, the newest 100 samples are shown. For more information about performance statistics, see [“About performance statistics” \(page 29\)](#).

You can update historical statistics.

To update displayed historical statistics

1. Display a historical statistics graph as described in [“Viewing performance statistics” \(page 135\)](#).
2. Select **Action > Update Historical Statistics**. The Update Historical Statistics panel opens and shows the default count value of 100.
3. To specify a count, in the Count field enter a value in the range of 5–100 and click **OK**.
4. To specify a time range, perform the following:
 - o Select the **Time Range** check box.
 - o Set date/time values for the starting and ending samples. The values must be between the current date/time and 6 months in the past. The ending values must be more recent than the starting values.

 **TIP:** If you specify a time range, it is recommended to specify a range of 24 hours or less.

- o Click **OK**.

In the Historical Data area of the Performance topic, the Time Range values are updated to show the times of the oldest and newest samples displayed. The graph for the selected components is updated.

Exporting historical performance statistics

You can export historical performance statistics in CSV format to a file on the network. You can then import the data into a spreadsheet or other third-party application.

The number of data samples downloaded is fixed at 100 to limit the size of the data file to be generated and transferred. The default is to retrieve all the available data (up to six months) aggregated into 100 samples. You can specify a different time range by specifying a start and end time. If the specified time range spans more than 100 15-minute samples, the data will be aggregated into 100 samples.


The resulting file will contain a row of property names and a row for each data sample.

To export historical performance statistics

1. In the Performance topic, from the Show list, select **Disks, Disk Groups, Virtual Pools, or Virtual Tiers**.
2. Select at least one component.

NOTE: Statistics are exported for all disks, regardless of which components are selected.

3. Select **Action > Export Historical Statistics**. The Export Historical Statistics panel opens.
4. To specify a time range, perform the following:
 - o Select the **Time Range** check box.
 - o Set date/time values for the starting and ending samples. The values must be between the current date/time and 6 months in the past. The ending values must be more recent than the starting values.

 **TIP:** If you specify a time range, it is recommended to specify a range of 24 hours or less.

5. Click **OK**.

NOTE: In Microsoft Internet Explorer, if the download is blocked by a security bar, select its **Download File** option. If the download does not succeed the first time, return to the Export Historical Statistics panel and retry the export operation.

6. When prompted to open or save the file, click **Save**.
 - o If you are using Firefox or Chrome and have a download directory set, the file `Disk_Performance.csv` is saved there.
 - o Otherwise, you are prompted to specify the file location and name. The default file name is `Disk_Performance.csv`. Change the name to identify the system, controller, and date.
7. Click **OK**.

Resetting performance statistics

You can reset (clear) the current or historical performance statistics for all components. When you reset statistics, an event is logged and new data samples will continue to be stored every five minutes.

To reset performance statistics

1. In the Performance topic, select **Action > Reset All Statistics**. The Reset All Statistics panel opens.
2. Perform one of the following:
 - o To reset current statistics, select **Current Data**.
 - o To reset historical statistics, select **Historical Data**.
3. Click **OK**. A confirmation panel appears.
4. Click **Yes** to continue. Otherwise, click **No**. If you clicked Yes, the statistics are cleared.

10 Working in the banner and footer



Banner and footer overview

The banner of the SMU interface contains four panels that are next to each other:

- The system panel shows system and firmware information.
- The connection information panel shows information about the link between the SMU and the storage system.
- The system date/time panel shows system date and time information.
- The user information panel shows the name of the logged-in user.

The footer of the SMU interface contains six panels that are next to each other:


- The system health panel shows the current health of the system and each controller.
- The event panel shows the last 1,000 or fewer events (organized by event type) that the system has logged.
- The capacity utilization panel shows a pair of color-coded bars that represent the physical capacity of the system and how the capacity is allocated and used.
- The host I/O panel shows a pair of color-coded bars for each controller that has active I/O, which represent the current IOPS for all ports and the current data throughput (MB/s) for all ports.
- The tier I/O panel shows a color-coded bar for each virtual pool (A, B, or both) that has active I/O.
- The I/O workload graph shows the relationship between the workload and the amount of storage capacity used.
- The activity panel shows notifications of recent system activities.

If you hover your cursor over any of these panels except for the activity panel, an additional panel with more detailed information displays. Some of these panels have menus that enable you to perform related tasks. There are two icons for panels that have a menu:  for the banner and  for the footer. Click anywhere in the panel to display the menu.

Viewing system information

The system panel in the banner shows the system name and the firmware bundle version installed for the controller that you are accessing.

Hover the cursor over this panel to display the System Information panel, which shows the system name, vendor, location, contact, and description. It also shows the firmware bundle version for each controller (A and B).

The  icon indicates that the panel has a menu. Click anywhere in the panel to display a menu to change system information settings ([page 52](#)) and system services settings ([page 51](#)), update firmware ([page 66](#)), restart or shut down controllers ([page 78](#)) and view SSL certificate information ([page 139](#)).

Viewing certificate information

By default, the system generates a unique SSL certificate for each controller. For the strongest security, you can replace the default system-generated certificate with a certificate issued from a trusted certificate authority.

The Certificate Information panel shows information for the active SSL certificates that are stored on the system for each controller. Tabs A and B 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 most often occurs when a controller is restarting, the certificate replacement process is still in process.

You can use your own certificates by uploading them through FTP or SFTP or by using the `contents` parameter of the `create certificate` CLI command to create certificates with your own unique certificate content. For a new certificate

to take effect, you must first restart the controller for it. For information on how to restart a controller, see “Restarting or shutting down controllers” (page 78).

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.




To view certificate information

1. In the banner, click the system panel and select **Show Certificate Info**. The Certificate Information panel opens.
2. After you have finished viewing certificate information, click **Close**.

Viewing connection information


The icon in the connection panel in the banner shows the current state of the management link between the SMU and the storage system. The connection information table shows the icon that displays for each state.

Table 15 Connection information

Icon	Meaning
	The management link is connected and the system is up. Animation shows when data is being transferred.
	The management link is connected but the system is down.
	The management link is not connected.

Hover the cursor over this panel to display the Connection Information panel, which shows the connection and system states.

Viewing system date and time information

The date/time panel in the banner shows the system date and time in the format *year-month-day hour:minutes:seconds*. The  icon indicates that the panel has a menu. Click anywhere in the panel to display a menu to change date and time settings.

Changing date and time settings

You can change the storage system date and time, which appear in the date/time panel in the banner. It is important to set the date and time so that entries in system logs and notifications have correct time stamps.

You can set the date and time manually or configure the system to use NTP to obtain them from a network-attached server. When NTP is enabled, and if an NTP server is available, the system time and date can be obtained from the NTP server. This allows multiple storage devices, hosts, log files, and so forth to be synchronized. If NTP is enabled but no NTP server is present, the date and time are maintained as if NTP was not enabled.

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

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

Whether NTP is enabled or disabled, the storage system does not automatically make time adjustments for Daylight Saving Time. You must make that adjustment manually.

To use manual date and time settings

1. In the banner, click the date/time panel and select **Set Date and Time**. The Set Date and Time panel opens.
2. Clear the **Network Time Protocol (NTP)** check box.
3. To set the Date value, enter the current date in the format *YYYY-MM-DD*.
4. To set the Time value, enter two-digit values for the hour and minutes and select either **AM**, **PM**, or **24H** (24-hour clock).
5. Click **OK**.


To obtain the date and time from an NTP server

1. In the banner, click the date/time panel and select **Set Date and Time**. The Set Date and Time panel opens.
2. Select the **Network Time Protocol (NTP)** check box.
3. Perform one of the following:
 - o To have the system retrieve time values from a specific NTP server, enter its IP address in the NTP Server Address field.
 - o To have the system listen for time messages sent by an NTP server in broadcast mode, clear the NTP Server Address field.
4. In the NTP Time Zone Offset field, enter the time zone as an offset in hours, and optionally, minutes, from UTC. For example, the Pacific Time Zone offset is -8 during Pacific Standard Time or -7 during Pacific Daylight Time. The offset for Bangalore, India is +5:30.
5. Click **OK**.

Viewing user information

The user panel in the banner shows the name of the signed-in user.


Hover the cursor over this panel to display the User Information panel, which shows the roles, accessible interfaces, and session timeout for this user.

The  icon indicates that the panel has a menu. Click anywhere in the panel to change settings for the signed-in user (monitor role) or to manage all users (manage role). For more information on user roles and settings, see [“Managing users” \(page 42\)](#).

Viewing health information

The health panel in the footer shows the current health of the system and each controller.

Hover the cursor over this panel to display the System Health panel, which shows the health state. If the system health is not OK, the System Health panel also shows information about resolving problems with unhealthy components.

The  icon indicates that the panel has a menu. Click anywhere in the panel to display a menu to change notification settings ([page 52](#)), save log data ([page 141](#)), and view system information ([page 61](#)).

Saving log data to a file

To help service personnel diagnose a system problem, you might be asked to provide system log data. Using the SMU, you can save the following log data to a compressed zip file:

- Device status summary, which includes basic status and configuration data for the system
- The event log from each controller
- The debug log from each controller
- The boot log, which shows the startup sequence, from each controller

- Critical error dumps from each controller, if critical errors have occurred
- CAPI traces from each controller

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

To save log data from the storage system to a network location

1. In the footer, click the health panel and select **Save Logs**. The Save Logs panel opens.
2. Enter your name, email address, and phone number so support personnel will know who provided the data. The contact name value can include a maximum of 100 bytes, using all characters except the following: “ ‘ ` & The email address can include a maximum of 100 characters., except the following: “ < > \ The phone number value can include only digits and no other characters.
3. Enter comments describing the problem and specifying the date and time when the problem occurred. This information helps service personnel when they analyze the log data. Comment text can include a maximum of 500 bytes.
4. Click **OK**. Log data is collected, which takes several minutes.






NOTE: In Microsoft Internet Explorer, if the download is blocked by a security bar, select its **Download File** option. If the download does not succeed the first time, return to the Save Logs panel and retry the save operation.

5. When prompted to open or save the file, click **Save**.
 - o If you are using Chrome, `store.zip` is saved to the downloads folder.
 - o If you are using Firefox and have a download folder set, `store.zip` is saved to that folder.
 - o Otherwise, you are prompted to specify the file location and name. The default file name is `store.zip`. Change the name to identify the system, controller, and date.

NOTE: The file must be uncompressed before the files it contains can be examined. The first file to examine for diagnostic data is `store_yyyy_mm_dd_hh_mm_ss.logs`.


Viewing event information

The event panel in the footer shows the numbers of the following types of events that the system has logged:

-  **Critical.** A failure occurred that may 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.

Hover the cursor over this area to display the Critical & Error Event Information panel, which shows:

- The number of events with Critical and Error severity that have occurred in the past 24 hours or in the last 1000 events
- The date and time when the last most-severe event occurred






The  icon indicates that the panel has a menu. Click anywhere in the panel to display a menu to view the most recent 1000 events on [“Viewing the event log” \(page 143\)](#) and change notification settings on [“Setting system notification settings” \(page 52\)](#).

Viewing the event log

If you are having a problem with the system, review the event log before calling technical support. Information shown in the event log might enable you to resolve the problem.

To view the event log, in the footer, click the events panel and select **Show Event List**. The Event Log Viewer panel opens. The panel shows a tabular view of the 1000 most recent events logged by either controller. All events are logged, regardless of notification settings. For information about notification settings, see [“Setting system notification settings” \(page 52\)](#). For information about using tables, see [“Tips for using tables” \(page 12\)](#). For information about events and the actions to take to solve them, see the Event Descriptions Reference Guide.

For each event, the panel shows the following information:

- Sev. One of the following severity icons:
 -  **Critical**. A failure occurred that may 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.
- Date/Time. The date and time when the event occurred, shown in the format *year-month-day hour:minutes:seconds*. Time stamps have one-second granularity.
- ID. The event ID. The prefix A or B identifies the controller that logged the event.
- Code. An event code that helps you and support personnel diagnose problems.
- Message. Brief information about the event. Click the message to show or hide additional information and recommended actions.
- Ctrl. The ID of the controller that logged the event.

When reviewing the event log, look for recent Critical, Error, or Warning events. For each, click the message to view additional information and recommended actions. Follow the recommended actions to resolve the problems.

Resources for diagnosing and resolving problems

- The troubleshooting chapter and LED descriptions appendix in your product's User Guide
- The topic about verifying component failure in the component's replacement instructions document
- The full list of event codes, descriptions, and recommended actions in your product's event documentation

Viewing capacity information

The capacity panel in the footer shows a pair of color-coded bars. The lower bar represents the physical capacity of the system and the upper bar identifies how the capacity is allocated and used. For color-code descriptions, see [“Color codes” \(page 14\)](#).

Hover the cursor over a segment to see the storage type and size represented by that segment. For instance, in a system where storage is being used, the bottom bar has color-coded segments that show the total unused disk space and space used by disk groups. The total of these segments is equal to the total disk capacity of the system.

In this same system, the top bar has color-coded segments for reserved, allocated, and unallocated space for virtual disk groups. If very little disk group space is used for any of these categories, it will not be visually represented.

Reserved space refers to space that is unavailable for host use. It consists of RAID parity and the metadata needed for internal management of data structures. Allocated space refers to the amount of space consumed by data written to the pool. Unallocated space is the difference between the space designated for all volumes and the allocated space.

Hover the cursor over a segment of a bar to see the storage size represented by that segment. Point anywhere in this panel to see the following information about capacity utilization in the Capacity Utilization panel:

- Total Disk Capacity. The total physical capacity of the system
- Unused. The total unused disk capacity of the system
- Global Spares. The total global spare capacity of the system
- Virtual Disk Groups. The capacity of the disk groups, both total and by pool
- Reserved. The reserved space for the disk groups, both total and by pool
- Allocated. The allocated space for the disk groups, both total and by pool
- Unallocated. The unallocated space for the disk groups, both total and by pool
- Uncommitted. For virtual disk groups, the uncommitted space in each pool (total space minus the allocated and unallocated space) and total uncommitted space

Viewing host I/O information

The host I/O panel in the footer shows a pair of color-coded bars for each controller that has active I/O. In each pair, the upper bar represents the current IOPS for all ports, which is calculated over the interval since these statistics were last requested or reset, and the lower bar represents the current data throughput (MB/s) for all ports, which is calculated over the interval since these statistics were last requested or reset. The pairs of bars are sized to represent the relative values for each controller. For color-code descriptions, see [“Color codes” \(page 14\)](#).

Hover the cursor over a bar to see the value represented by that bar.

Hover the cursor anywhere in the panel to display the Host I/O Information panel, which shows the current port IOPS and data throughput (MB/s) values for each controller.

Viewing tier I/O information

The tier I/O panel in the footer shows a color-coded bar for each virtual pool (A, B, or both) that has active I/O. The bars are sized to represent the relative IOPS for each pool. Each bar contains a segment for each tier that has active I/O. The segments are sized to represent the relative IOPS for each tier. For color-code descriptions, see [“Color codes” \(page 14\)](#).

Hover the cursor over a segment to see the value represented by that segment.

Hover the cursor anywhere in this panel to display the Tier I/O Information panel, which shows the following details for each tier in each virtual pool:

- Current IOPS for the pool, calculated over the interval since these statistics were last requested or reset.
- Current data throughput (MB/s) for the pool, calculated over the interval since these statistics were last requested or reset.

The panel also contains combined total percentages of IOPS and current data throughput (MB/s) for both pools.

Viewing I/O workload activity

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

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

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

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

You can set the following options:

- **Pools**—Select whether to calculate the workload for either pool A or B.
- **Values**—Select whether to base the calculations on either 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. There are multiple samples 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, instead. This leaves one value for each bin whether average or peak was selected. From there, workload calculations are made using the bin numbers as input.
- **Show**—Select whether to limit the data used for calculations to small read I/Os only, small write I/Os only, or the combined total of small read and write I/Os. Small I/Os are random access operations, as opposed to large I/Os which are sequential access operations.
- **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 Current SSD Capacity and total Pool Capacity display above 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 when making a purchase.
- 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 SSD purchase.

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

Viewing recent system activity

The activity panel in the footer shows notifications of recent system activities, such as the loading of configuration data upon sign-in, events with the Resolved status, and scheduled tasks.

Viewing the notification history

The Notification History panel shows past activity notifications for this SMU session. You can page through listed items by using the following buttons:

- Show next set of items.
- Reached end of list.
- Show previous set of items.
- *Reached start of list.

When you sign out, the list is cleared.

To view notification history

1. Click the activity panel in the footer and select **Notification History**. The Notification History panel opens.
2. View activity notifications, using the navigation buttons.
3. Click **Close** when you are finished.

11 Support and other resources

Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:
www.hpe.com/assistance
- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:
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, go to either of the following:

Hewlett Packard Enterprise Support Center

www.hpe.com/support/hpesc

Hewlett Packard Enterprise Support Center: Software downloads

www.hpe.com/support/downloads

Software Depot

www.hpe.com/support/softwaredepot

- To subscribe to eNewsletters and alerts:
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 HP Support Materials** page:
www.hpe.com/support/AccessToSupportMaterials

① **IMPORTANT:** Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Passport set up with relevant entitlements.

Customer self repair

Hewlett Packard Enterprise customer self repair (CSR) programs allow you to repair your product. If a CSR part needs to be replaced, it will be shipped directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your Hewlett Packard Enterprise authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider or go to the CSR website:

www.hpe.com/support/selfrepair

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 will initiate a fast and accurate resolution based on your product's service level. 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.

Remote support and Proactive Care information

HPE Get Connected

www.hpe.com/services/getconnected

HPE Proactive Care services

www.hpe.com/services/proactivecare

HPE Proactive Care service: Supported products list

www.hpe.com/services/proactivecaresupportedproducts

HPE Proactive Care advanced service: Supported products list

www.hpe.com/services/proactivecareadvancedsupportedproducts

Proactive Care customer information

Proactive Care central

www.hpe.com/services/proactivecarecentral

Proactive Care service activation

www.hpe.com/services/proactivecarecentralgetstarted

Warranty information

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

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

Additional warranty information

HPE ProLiant and x86 Servers and Options

www.hpe.com/support/ProLiantServers-Warranties

HPE Enterprise Servers

www.hpe.com/support/EnterpriseServers-Warranties

HPE Storage Products

www.hpe.com/support/Storage-Warranties

HPE Networking Products

www.hpe.com/support/Networking-Warranties

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:

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:

www.hpe.com/info/reach

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

www.hpe.com/info/ecodata

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

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, send any errors, suggestions, or comments to Documentation Feedback (docsfeedback@hpe.com). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.

A Other management interfaces

SNMP reference

This appendix describes the Simple Network Management Protocol (SNMP) capabilities that HPE MSA 1050/2050 storage systems support. This includes standard MIB-II, the FibreAlliance SNMP Management Information Base (MIB) version 2.2 objects, and enterprise traps.

The storage systems can report their status through SNMP. SNMP provides basic discovery using MIB-II, more detailed status with the FA MIB 2.2, and asynchronous notification using enterprise traps.

SNMP is a widely used network monitoring and control protocol. It is an application layer protocol that facilitates the exchange of management information between network devices. It is part of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite.

SNMP enables network administrators to manage network performance, find and solve network problems, and plan for network growth. Data is passed from SNMP agents reporting activity on each network device to the workstation console used to oversee the network. The agents return information contained in a Management Information Base (MIB), which is a data structure that defines what is obtainable from the device and what can be controlled (turned on and off, etc.).

Supported SNMP versions

The storage systems allow use of SNMPv2c or SNMPv3. SNMPv2c uses a community-based security scheme. For improved security, SNMPv3 provides authentication of the network management system that is accessing the storage system, and encryption of the information transferred between the storage system and the network management system.

When SNMPv3 is disabled, SNMPv2c will be active. When SNMPv3 is enabled, SNMPv2c will only have access to the MIB-II common system information. This allows device discovery.

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 1050/2050 storage systems is 1.3.6.1.4.111.2.51, 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 occurring in the storage system. These events can be selected by severity and by individual event type. A maximum of three SNMP trap destinations can be configured by IP address.

Enterprise event severities are informational, minor, major, and critical. There is a different trap type for each of these severities. The trap format is represented by the enterprise traps MIB, `msa2000traps.mib`. Information included is the event ID, the event code type, and a text description generated from the internal event. Equivalent information can also be sent using email or popup alerts to users who are logged in to the SMU.

The text of the trap MIB is included at the end of this appendix.

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 and 4.0 differences” \(page 162\)](#).

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

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

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

The port section is only relevant to products with Fibre Channel host ports.

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

The traps section is not supported. It has been replaced by an ability to configure trap destinations using the CLI or the SMU. The statistics section is not implemented.

The following table lists the MIB objects, their descriptions and the value set in MSA 1050/2050 storage systems. Unless specified otherwise, objects are *not* settable.

Table 16 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, <code>http://10.1.2.3</code> . If a web server is not present on the device, this string is empty in accordance with the FA MIB2.2 Spec.	Default: <code>http://10.0.0.1</code>
StatusChangeTime	<code>sysuptime</code> timestamp of the last status change event, in centiseconds. <code>sysuptime</code> starts at 0 when the Storage Controller boots and keeps track of the up time. <code>statusChangeTime</code> is updated each time an event occurs.	0 at startup
ConfigurationChangeTime	<code>sysuptime</code> timestamp of the last configuration change event, in centiseconds. <code>sysuptime</code> starts at 0 when the Storage Controller boots and keeps track of the up time. <code>configurationChangeTime</code> is updated each time an event occurs.	0 at startup
ConnUnitTableChangeTime	<code>sysuptime</code> timestamp of the last update to the <code>connUnitTable</code> (an entry was either added or deleted), in centiseconds	0 always (entries are not added to or deleted from the <code>connUnitTable</code>)

Table 16 FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
connUnitTable	Includes the following objects as specified by the FA MIB2.2 Spec	
connUnitId	Unique identification for this connectivity unit	Total of 16 bytes comprised of 8 bytes of the node WWN or similar serial number-based identifier (for example, 1000005013b05211) with the trailing 8 bytes equal to zero
connUnitGlobalId	Same as connUnitId	Same as connUnitId
connUnitType	Type of connectivity unit	storage-subsystem(11)
connUnitNumports	Number of host ports in the connectivity unit	Number of host ports
connUnitState	Overall state of the connectivity unit	online(2) or unknown(1), as appropriate
connUnitStatus	Overall status of the connectivity unit	ok(3), warning(4), failed(5), or unknown(1), as appropriate
connUnitProduct	Connectivity unit vendor's product model name	Model string
connUnitSn	Serial number for this connectivity unit	Serial number string
connUnitUpTime	Number of centiseconds since the last unit initialization	0 at startup
connUnitUrl	Same as systemURL	Same as systemURL
connUnitDomainId	Not used; set to all 1s as specified by the FA MIB2.2 Spec	0xFFFF
connUnitProxyMaster	Stand-alone unit returns yes for this object	yes(3) since this is a stand-alone unit
connUnitPrincipal	Whether this connectivity unit is the principal unit within the group of fabric elements. If this value is not applicable, returns unknown.	unknown(1)
connUnitNumSensors	Number of sensors in the connUnitSensorTable	33
connUnitStatusChangeTime	Same as statusChangeTime	Same as statusChangeTime
connUnitConfigurationChangeTime	Same as configurationChangeTime	Same as configurationChangeTime
connUnitNumRevs	Number of revisions in the connUnitRevsTable	16
connUnitNumZones	Not supported	0
connUnitModuleId	Not supported	16 bytes of 0s
connUnitName	Settable: Display string containing a name for this connectivity unit	Default: Uninitialized Name
connUnitInfo	Settable: Display string containing information about this connectivity unit	Default: Uninitialized Info

Table 16 FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
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
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” (page 156)
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” (page 156)
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” (page 158)
connUnitSensorName	Display string containing textual identification of the sensor intended primarily for operator use	See “External details for connUnitSensorTable” (page 158)
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

Table 16 FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
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” (page 158)
connUnitSensorCharacteristic	Characteristics being monitored by this sensor	See “External details for connUnitSensorTable” (page 158)
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
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)
connUnitPortTransmitterType	Technology of the port transceiver	unknown(1) for Fibre Channel ports
connUnitPortModuleType	Module type of the port connector	unknown(1)
connUnitPortWwn	Fibre Channel World Wide Name (WWN) of the port if applicable	WWN octet for the port, or empty string if the port is not present
connUnitPortFCId	Assigned Fibre Channel ID of this port	Fibre Channel ID of the port All bits set to 1 if the Fibre Channel ID is not assigned or if the port is not present
connUnitPortSn	Serial number of the unit (for example, for a GBIC). If this is not applicable, returns an empty string.	Empty string
connUnitPortRevision	Port revision (for example, for a GBIC)	Empty string

Table 16 FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
connUnitPortVendor	Port vendor (for example, for a GBIC)	Empty string
connUnitPortSpeed	Speed of the port in KByte per second (1 KByte = 1000 Byte)	Port speed in KByte per second, or 0 if the port is not present
connUnitPortControl	Not supported	invalid(2) for an SNMP GET operation and not settable through an SNMP SET operation
connUnitPortName	String describing the addressed port	See “ External details for connUnitPortTable ” (page 159)
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
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

Table 16 FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
trapRegTable	Includes the following objects per the FA MIB2.2 Spec	
trapRegIpAddress	IP address of a client registered for traps	IP address set by user
trapRegPort	User Datagram Protocol (UDP) port to send traps to for this host	162
trapRegFilter	Settable: Defines the trap severity filter for this trap host. The connUnit will send traps to this host that have a severity level less than or equal to this value.	Default: warning(6)
trapRegRowState	Specifies the state of the row	READ: rowActive(3) if traps are enabled. Otherwise rowInactive(2) WRITE: Not supported
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 2050
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 section specify values for certain objects described in [Table 16](#).

External details for connUnitRevsTable

Table 17 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)

Table 17 connUnitRevsTable index and description values (continued)

connUnitRevsIndex	connUnitRevsDescription
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 18 connUnitSensorTable index, name, type, and characteristic values

connUnitSensorIndex	connUnitSensorName	connUnitSensorType	connUnitSensorCharacteristic
1	Onboard Temperature 1 (Controller A)	board(8)	temperature(3)
2	Onboard Temperature 1 (Controller B)	board(8)	temperature(3)
3	Onboard Temperature 2 (Controller A)	board(8)	temperature(3)
4	Onboard Temperature 2 (Controller B)	board(8)	temperature(3)
5	Onboard Temperature 3 (Controller A)	board(8)	temperature(3)
6	Onboard Temperature 3 (Controller B)	board(8)	temperature(3)
7	Disk Controller Temperature (Controller A)	board(8)	temperature(3)
8	Disk Controller Temperature (Controller B)	board(8)	temperature(3)
9	Memory Controller Temperature (Controller A)	board(8)	temperature(3)
10	Memory Controller Temperature (Controller B)	board(8)	temperature(3)
11	Capacitor Pack Voltage (Controller A)	board(8)	power(9)
12	Capacitor Pack Voltage (Controller B)	board(8)	power(9)
13	Capacitor Cell 1 Voltage (Controller A)	board(8)	power(9)
14	Capacitor Cell 1 Voltage (Controller B)	board(8)	power(9)
15	Capacitor Cell 2 Voltage (Controller A)	board(8)	power(9)
16	Capacitor Cell 2 Voltage (Controller B)	board(8)	power(9)
17	Capacitor Cell 3 Voltage (Controller A)	board(8)	power(9)
18	Capacitor Cell 3 Voltage (Controller B)	board(8)	power(9)
19	Capacitor Cell 4 Voltage (Controller A)	board(8)	power(9)
20	Capacitor Cell 4 Voltage (Controller B)	board(8)	power(9)
21	Capacitor Charge Percent (Controller A)	board(8)	other(2)
22	Capacitor Charge Percent (Controller B)	board(8)	other(2)
23	Overall Status	enclosure(7)	other(2)
24	Upper IOM Temperature (Controller A)	enclosure(7)	temperature(3)
25	Lower IOM Temperature (Controller B)	enclosure(7)	temperature(3)
26	Power Supply 1 (Left) Temperature	power-supply(5)	temperature(3)
27	Power Supply 2 (Right) Temperature	power-supply(5)	temperature(3)
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 18 connUnitSensorTable index, name, type, and characteristic values (continued)

connUnitSensorIndex	connUnitSensorName	connUnitSensorType	connUnitSensorCharacteristic
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 19 connUnitPortTable index and name values

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

Configuring SNMP event notification in the SMU

1. Verify that the storage system's SNMP service is enabled. [“Enabling or disabling system-management services” \(page 50\)](#).
2. Configure and enable SNMP traps.
3. Optionally, configure a user account to receive SNMP traps.

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. See [“Enabling or disabling system-management services” \(page 50\)](#). 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 [“Managing users” \(page 42\)](#). The same users, security protocols, and passwords must be configured in the network management system.

Enterprise trap MIB

The following pages show the source for the msa2000traps.mib. This MIB defines the content of the SNMP traps that HPE MSA 1050/2050 storage systems generate.

```
-----
-- MSA2000 Array MIB for SNMP Traps
--
-- $Revision: 11692 $
--
-- Copyright (c) 2008 Hewlett-Packard Development Company, L.P.
-- Copyright (c) 2005-2008 Dot Hill Systems Corp.
-- Confidential computer software. Valid license from HP required for possession,
-- use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer
-- Software, Computer Software Documentation, and Technical Data for Commercial
-- Items are licensed to the U.S. Government under vendor's standard commercial
-- license.
--
--   MSA2000traps MIB Revision
--   =====
-- Revision 1.1  2008/02/27
-- Initial revision
-- Revision 1.2  2008/03/18
-- Updated copyright notice
--
-----

MSA2000TRAPS-MIB
-- Last edit date: Feb 27th, 2008
DEFINITIONS ::= BEGIN
    IMPORTS
        enterprises
            FROM RFC1155-SMI
        TRAP-TYPE
            FROM RFC-1215
        connUnitEventId, connUnitEventType, connUnitEventDescr
            FROM FA-MIB40;

    --Textual conventions for this MIB

-----

    -- vendor
    hp      OBJECT IDENTIFIER ::= { enterprises 11 }
    nm      OBJECT IDENTIFIER ::= { hp 2 }
    hpMSA   OBJECT IDENTIFIER ::= { nm 51 }

-- Related traps

    msaEventInfoTrap TRAP-TYPE
```



```

ENTERPRISE hpMSA
VARIABLES { connUnitEventId,
             connUnitEventType,
             connUnitEventDescr }
DESCRIPTION
    "An event has been generated by the storage array.
    Recommended severity level (for filtering): info"

-- Trap annotations are as follows:
--#TYPE "Informational storage event"
--#SUMMARY "Informational storage event # %d, type %d, description: %s"
--#ARGUMENTS {0,1,2}
--#SEVERITY INFORMATIONAL
--#TIMEINDEX 6
::= 3001

msaEventWarningTrap TRAP-TYPE
ENTERPRISE hpMSA
VARIABLES { connUnitEventId,
             connUnitEventType,
             connUnitEventDescr }
DESCRIPTION
    "An event has been generated by the storage array.
    Recommended severity level (for filtering): warning"

-- Trap annotations are as follows:
--#TYPE "Warning storage event"
--#SUMMARY "Warning storage event # %d, type %d, description: %s"
--#ARGUMENTS {0,1,2}
--#SEVERITY MINOR
--#TIMEINDEX 6
::= 3002

msaEventErrorTrap TRAP-TYPE
ENTERPRISE hpMSA
VARIABLES { connUnitEventId,
             connUnitEventType,
             connUnitEventDescr }
DESCRIPTION
    "An event has been generated by the storage array.
    Recommended severity level (for filtering): error"

-- Trap annotations are as follows:
--#TYPE "Error storage event"
--#SUMMARY "Error storage event # %d, type %d, description: %s"
--#ARGUMENTS {0,1,2}
--#SEVERITY MAJOR
--#TIMEINDEX 6
::= 3003

msaEventCriticalTrap TRAP-TYPE

```

```

ENTERPRISE hpMSA
VARIABLES { connUnitEventId,
             connUnitEventType,
             connUnitEventDescr }
DESCRIPTION
    "An event has been generated by the storage array.
    Recommended severity level (for filtering): critical"

-- Trap annotations are as follows:
--#TYPE "Critical storage event"
--#SUMMARY "Critical storage event # %d, type %d, description: %s"
--#ARGUMENTS {0,1,2}
--#SEVERITY CRITICAL
--#TIMEINDEX 6
 ::= 3004

msaEventResolvedTrap TRAP-TYPE
ENTERPRISE hpMSA
VARIABLES { connUnitEventId,
             connUnitEventType,
             connUnitEventDescr }
DESCRIPTION
    "An issue has been resolved on the array.
    Recommended severity level (for filtering): resolved"

-- Trap annotations are as follows:
--#TYPE "Resolved storage event"
--#SUMMARY "Resolved storage event # %d, type %d, description: %s"
--#ARGUMENTS {0,1,2}
--#SEVERITY INFORMATIONAL
--#TIMEINDEX 6
 ::= 3005

END

```

FA MIB 2.2 and 4.0 differences

FA MIB 2.2 is a subset of FA MIB 4.0. Therefore, SNMP elements implemented in MSA 1050/2050 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 FTP and SFTP

Although the SMU is the preferred interface for downloading log data and historical disk-performance statistics, updating firmware, installing a license, you can also use FTP and SFTP to do these tasks — and to install security certificates and keys.

-
- ❗ **IMPORTANT:** Do not attempt to do more than one of the operations in this appendix at the same time. They can interfere with each other and the operations may fail. Specifically, do not try to do more than one firmware update at the same time or try to download system logs while doing a firmware update.
-

Downloading system logs

To help service personnel diagnose a system problem, you might be asked to provide system log data. You can download this data by accessing the system's FTP or SFTP interface and running the `get logs` command. When both controllers are online, regardless of operating mode, `get logs` will download a single, compressed zip file that includes:

- Device status summary, which includes basic status and configuration data for the system
- 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 FTP/SFTP client. A GUI-based FTP/SFTP client might not work.

To download system logs

1. In the SMU, prepare to use FTP/SFTP:
 - a. Determine the network-port IP addresses of the system's controllers. See [“Configuring controller network ports” \(page 49\)](#).
 - b. Verify that the system's FTP/SFTP service is enabled and take note of the FTP/SFTP service port. See [“Enabling or disabling system-management services” \(page 50\)](#).
 - c. Verify that the user you will log in as has permission to use the FTP interface. The same setting allows a user to transfer files using both FTP and SFTP.
2. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the destination directory for the log file.
3. Using the FTP/SFTP port specified in the system services settings, enter:


```
psftp controller-network-address -P port or ftp controller-network-address
```

 For example:


```
psftp 10.235.216.152 -P 1022
ftp 10.1.0.9
```
4. Log in as a user that has permission to use the FTP/SFTP 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
```

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.

6. Quit the FTP/SFTP session.

NOTE: You must uncompress a zip file before you can view the files it contains. To examine diagnostic data, first view `store_yyyy_mm_dd_hh_mm_ss.logs`.

Transferring log data to a log-collection system

If the log-management feature is configured in pull mode, a log-collection system can access the storage system's FTP or SFTP interface and use the `get managed-logs` command to retrieve untransferred data from a system log file. This command retrieves the untransferred data from the specified log to a compressed zip file on the log-collection system. Following the transfer of a log's data, the log's capacity status is reset to zero indicate that there is no untransferred data. Log data is controller specific.

For an overview of the log-management feature, see [“About managed logs” \(page 30\)](#).

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

To transfer log data to a log-collection system

1. In the SMU, prepare to use FTP/SFTP:

- a. Determine the network-port IP addresses of the system's controllers. See [“Configuring controller network ports” \(page 49\)](#).
- b. Verify that the system's FTP/SFTP service is enabled. See [“Enabling or disabling system-management services” \(page 50\)](#).
- c. Verify that the user you will log in as has permission to use the FTP/SFTP interface. The same setting allows a user to transfer files using both FTP and SFTP. See [“To modify a user” \(page 45\)](#).

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:

```
psftp controller-network-address -P port or ftp controller-network-address
```

For example:

```
psftp 10.235.216.152 -P 1022
```

```
ftp 10.1.0.9
```

4. Log in as a user that has permission to use the FTP/SFTP interface.

5. Enter:

```
get managed-logs:log-type filename.zip
```

where:

- o *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.

- o *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 FTP/SFTP 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 FTP/SFTP 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 six months) aggregated into 100 samples. You can specify a different time range by specifying a start and end time. If the specified time range spans more than 100 15-minute samples, the data will be aggregated into 100 samples.

The resulting file will contain a row of property names and a row for each data sample, as shown in the following example. For property descriptions, see the topic about the `disk-hist-statistics` basetype in the CLI Reference Guide.

```
"sample-time", "durable-id", "serial-number", "number-of-ios", ...
"2012-01-26 01:00:00", "disk_1.1", "PLV2W1XE", "2467917", ...
"2012-01-26 01:15:00", "disk_1.1", "PLV2W1XE", "2360042", ...
...
```

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

To retrieve historical disk-performance statistics

1. In the SMU, prepare to use FTP/SFTP:
 - a. Determine the network-port IP addresses of the system's controllers. See [“Configuring controller network ports” \(page 49\)](#).
 - b. Verify that the system's FTP/SFTP service is enabled. See [“Enabling or disabling system-management services” \(page 50\)](#).
 - c. Verify that the user you will log in as has permission to use the FTP/SFTP interface. The same setting allows a user to transfer files using both FTP and SFTP. See [“To modify a user” \(page 45\)](#).

2. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the destination directory for the log file.

3. Enter:

```
psftp controller-network-address -P port or ftp controller-network-address
```

For example:

```
psftp 10.235.216.152 -P 1022
```

```
ftp 10.1.0.9
```

4. Log in as a user that has permission to use the FTP/SFTP interface.

5. Enter:

```
get perf[:date/time-range] filename.csv
```

where:

- o *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.
- o *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
```


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 FTP/SFTP session.

Updating firmware

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

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

 **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:**

- 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 on the System Overview panel ([page 141](#)).
-

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 disk-drive firmware by searching on <http://www.hpe.com/storage/msadrivefirmware> for your drive model number to find the latest firmware to download. Then, load the firmware file obtained from the HPE web download site at <http://www.hpe.com/support/hpesc>. To install an HPE ROM Flash Component or firmware Smart Component, follow the instructions on the HPE web site. 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. As a user with a `manage` role, obtain the appropriate firmware file and download it to your computer or network.
2. In the SMU, prepare to use FTP/SFTP:
 - a. Determine the network-port IP addresses of the system's controllers.
 - b. Verify that the system's FTP/SFTP service is enabled.
 - c. Verify that the user you will log in as has permission to use the FTP/SFTP interface. The same setting allows a user to transfer files using both FTP and SFTP.
3. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the directory containing the firmware file to load.

4. Enter:

```
psftp controller-network-address -P port or ftp controller-network-address
```

For example:

```
psftp 10.235.216.152 -P 1022
```

```
ftp 10.1.0.9
```

5. Log in as an FTP/SFTP 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 FTP/SFTP prompt is redisplayed. The code is not loaded.

Firmware update typically takes 10 minutes for a controller having current CPLD firmware, or 20 minutes for a controller with downlevel CPLD firmware. If the controller enclosure has attached enclosures, allow additional time for each expansion module's enclosure management processor (EMP) to be updated. This typically takes 2.5 minutes for each EMP in an MSA 1050/2050 drive enclosure.

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

If the Storage Controller cannot be updated, the update operation is canceled. If the FTP/SFTP prompt does not return, quit the FTP/SFTP session and log in again. Verify that you specified the correct firmware file and repeat the update. If this problem persists, contact technical support.

When firmware update on the local controller is complete, the FTP/SFTP session returns to the `sftp>` prompt, and the FTP/SFTP 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, allow an additional 5–20 minutes for both controllers to be updated.

7. Quit the FTP/SFTP 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

A drive enclosure can contain one or two expansion modules. Each expansion module contains an enclosure management processor (EMP). All modules of the same product model should run the same firmware version.

You can update disk-drive firmware by searching on <http://www.hpe.com/storage/msadrivefirmware> for your drive model number to find the latest firmware to download. Then, load the firmware file obtained from the HPE web download site at <http://www.hpe.com/support/hpesc>. To install an HPE ROM Flash Component or firmware Smart Component, follow the instructions on the HPE web site. Otherwise, to install a firmware binary file, follow the steps below.

You can specify to update all expansion modules or only specific expansion modules. If you specify to update all expansion modules and the system contains more than one type of enclosure, the update will be attempted on all enclosures in the system. The update will only succeed for enclosures whose type matches the file, and will fail for enclosures of other types.

To update expansion-module firmware

1. As a user with a `manage` role, obtain the appropriate firmware file and download it to your computer or network.
2. If you want to update all expansion modules, continue with the next step. Otherwise, in the SMU, determine the address of each expansion module to update:
 - a. In the Configuration View panel, select a drive enclosure.
 - b. In the enclosure properties table, note each EMP's bus ID and target ID values. For example, 0 and 63, and 1 and 63. Bus 0 is the bus that is native to a given controller, while bus 1 is an alternate path through the partner controller. It is recommended to perform update tasks consistently through one controller to avoid confusion.
3. In the SMU, prepare to use FTP/SFTP:
 - a. Determine the network-port IP addresses of the system's controllers.
 - b. Verify that the system's FTP/SFTP service is enabled.
 - c. Verify that the user you will log in as has permission to use the FTP/SFTP interface. The same setting allows a user to transfer files using both FTP and SFTP.
4. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the directory containing the firmware file to load.
5. Enter:

```
psftp controller-network-address -P port or ftp controller-network-address
```

For example:

```
psftp 10.235.216.152 -P 1022
```

```
ftp 10.1.0.9
```
6. Log in as an FTP/SFTP user.

7. Either:

- o To update all expansion modules, enter:
`put firmware-file encl`
- o To update specific expansion modules, enter:
`put firmware-file encl:EMP-bus-ID:EMP-target-ID`

△ CAUTION: Do not perform a power cycle or controller restart during the firmware update. If the update is interrupted or there is a power failure, the module might become inoperative. If this occurs, contact technical support. The module might need to be returned to the factory for reprogramming.

It typically takes 2.5 minutes to update each EMP in an MSA 1050/2050 drive enclosure. In FTP, wait for a message that the code load has completed. No messages are displayed in SFTP.

NOTE: If the update fails, verify that you specified the correct firmware file and try the update a second time. If it fails again, contact technical support.

8. If you are updating specific expansion modules, repeat [step 7](#) for each remaining expansion module that needs to be updated.
9. Quit the FTP/SFTP session.
10. Verify that each updated expansion module has the correct firmware version.

Updating disk firmware

You can update disk-drive firmware by searching on <http://www.hpe.com/storage/msadrivefirmware> for your drive model number to find the latest firmware to download. Then, load the firmware file obtained from the HPE web download site at <http://www.hpe.com/support/hpesc>. To install an HPE ROM Flash Component or firmware Smart Component, follow the instructions on the HPE web site. Otherwise, to install a firmware binary file, follow the steps below.

A dual-ported disk can be updated from either controller.

NOTE: Disks of the same model in the storage system must have the same firmware revision.

You can specify to update all disks or only specific disks. If you specify to update all disks and the system contains more than one type of disk, the update will be attempted on all disks in the system. The update will only succeed for disks whose type matches the file, and will fail for disks of other types.

To prepare for update

1. As a user with a manage role, obtain the appropriate firmware file and download it to your computer or network.
2. Check the disk manufacturer's documentation to determine whether disks must be power cycled after firmware update.
3. If you want to update all disks of the type that the firmware applies to, continue with the next step. Otherwise, in the SMU, for each disk to update:
 - a. Determine the enclosure number and slot number of the disk.
 - b. If the disk is associated with a disk group and is single ported, determine which controller owns the disk group.

4. In the SMU, prepare to use FTP/SFTP:
 - a. Determine the network-port IP addresses of the system's controllers.
 - b. Verify that the system's FTP/SFTP service is enabled.
 - c. Verify that the user you will log in as has permission to use the FTP interface. The same setting allows a user to transfer files using both FTP and SFTP.
5. Stop I/O to the storage system. During the update all volumes will be temporarily inaccessible to hosts. If I/O is not stopped, mapped hosts will report I/O errors. Volume access is restored after the update completes.

To update disk firmware

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

```
psftp controller-network-address -P port or ftp controller-network-address
```

For example:

```
psftp 10.235.216.152 -P 1022
ftp 10.1.0.9
```
3. Log in as an FTP/SFTP user.
4. Either:
 - o To update all disks of the type that the firmware applies to, enter:

```
put firmware-file disk
```
 - o To update specific disks, enter:

```
put firmware-file disk:enclosure-ID:slot-number
```

For example:

```
put firmware-file disk:1:11
```

△ CAUTION: Do not power cycle enclosures or restart a controller during the firmware update. If the update is interrupted or there is a power failure, the disk might become inoperative. If this occurs, contact technical support.

It typically takes several minutes for the firmware to load. In FTP, wait for the message `Operation Complete` to appear. No messages are displayed in SFTP.

NOTE: If the update fails, verify that you specified the correct firmware file and try the update a second time. If it fails again, contact technical support.

5. If you are updating specific disks, repeat [step 4](#) for each remaining disk to update.
6. Quit the FTP/SFTP session.
7. If the updated disks must be power cycled:
 - a. Shut down both controllers by using the SMU.
 - b. Power cycle all enclosures as described in your product's User Guide.
8. Verify that each disk has the correct firmware revision.

Installing a license file

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

```
psftp controller-network-address -P port or ftp controller-network-address
```

For example:

```
psftp 10.235.216.152 -P 1022
```

```
ftp 10.1.0.9
```

5. Log in as an FTP/SFTP 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.

As an alternative to using the CLI to create a security certificate on the storage system, you can use FTP/SFTP to install a custom certificate on the system. A certificate consists of a certificate file and an associated key file. The certificate can be created by using OpenSSL, for example, and is expected to be valid. If you replace the controller module in which a custom certificate is installed, the partner controller will automatically install the certificate file to the replacement controller module.

To install a security certificate

1. In the SMU, prepare to use FTP/SFTP:
 - a. Determine the network-port IP addresses of the system's controllers. See [“Configuring controller network ports” \(page 49\)](#).
 - b. Verify that the system's FTP/SFTP service is enabled. See [“Enabling or disabling system-management services” \(page 50\)](#).
 - c. Verify that the user you will log in as has permission to use the FTP/SFTP interface. The same setting allows a user to transfer files using both FTP and SFTP. See [“To modify a user” \(page 45\)](#).
2. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the directory that contains the certificate files.
3. Enter:

```
psftp controller-network-address -P port or ftp controller-network-address
```

For example:

```
psftp 10.235.216.152 -P 1022
```

```
ftp 10.1.0.9
```

4. Log in as a user that has permission to use the FTP/SFTP interface.
5. Enter:


```
put certificate-file-name cert-file
```

 where *certificate-file-name* is the name of the certificate file for your specific system.
6. Enter:


```
put key-file-name cert-key-file
```

 where *key-file-name* is the name of the security key file for your specific system.
7. Restart both Management Controllers to have the new security certificate 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 FTP/SFTP interface and use the `get logs` command with the `heatmap` option to download a log file in CSV format. The file contains data for the past seven days from both controllers.

1. In the SMU, prepare to use FTP/SFTP:
 - a. Determine the network-port IP addresses of the system's controllers. See ["Configuring controller network ports" \(page 49\)](#).
 - b. Verify that the system's FTP/SFTP service is enabled. See ["Enabling or disabling system-management services" \(page 50\)](#).
 - c. Verify that the user you will log in as has permission to use the FTP/SFTP interface. The same setting allows a user to transfer files using both FTP and SFTP. See ["To modify a user" \(page 45\)](#).
2. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the destination directory for the log file.
3. Enter:


```
psftp controller-network-address -P port or ftp controller-network-address
```

 For example:


```
psftp 10.235.216.152 -P 1022
ftp 10.1.0.9
```
4. Log in as a user that has permission to use the FTP/SFTP 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 FTP/SFTP session.

Using SMI-S

This appendix provides information for network administrators who are managing the storage system from a storage management application through the Storage Management Initiative Specification (SMI-S). SMI-S is a Storage Networking Industry Association (SNIA) standard that enables interoperable management for storage networks and storage devices.

SMI-S replaces multiple disparate managed object models, protocols, and transports with a single object-oriented model for each type of component in a storage network. The specification was created by SNIA to standardize storage management solutions. SMI-S enables management applications to support storage devices from multiple vendors

quickly and reliably because they are no longer proprietary. SMI-S detects and manages storage elements by type, not by vendor.

The key SMI-S components are:

- Web-based Enterprise Management (WBEM). A set of management and internet standard technologies developed to unify the management of enterprise computing environments. WBEM includes the following specifications:
 - CIM XML: defines XML elements, conforming to DTD, which can be used to represent CIM classes and instances
 - CIMxml Operations over HTTP/HTTPS: defines a mapping of CIM operations onto HTTP/HTTPS; used as a transport mechanism
- Common Information Model (CIM). The data model for WBEM. Provides a common definition of management information for systems, networks, applications and services, and allows for vendor extensions. SMI-S is the interpretation of CIM for storage. It provides a consistent definition and structure of data, using object-oriented techniques. The standard language used to define elements of CIM is MOF.
- Service Location Protocol (SLP). Enables computers and other devices to find services in a local area network without prior configuration. SLP has been designed to scale from small, unmanaged networks to large enterprise networks.

Embedded SMI-S array provider

The embedded SMI-S array provider provides an implementation of SMI-S 1.5 using `cim-xml` over HTTP/HTTPS. SMI-enabled management clients such as HPE SIM or HPE Storage Essentials can perform storage management tasks such as monitoring, configuring or event-management. The provider supports the Array and Server profiles with additional (or supporting) subprofiles. The Server profile provides a mechanism to tell the client how to connect and use the embedded provider. The Array profile has the following supporting profiles and subprofiles:

- Array profile
- Block Services package
- Physical Package package
- Health package
- Multiple Computer System subprofile
- Masking and Mapping profile
- FC Initiator Ports profile
- SAS Initiator Ports profile
- iSCSI Initiator Ports profile
- Disk Drive Lite profile
- Extent Composition subprofile
- Storage Enclosure profile
- Fan profile
- Power Supply profile
- Sensors profile
- Access Points subprofile
- Location subprofile
- Software Inventory subprofile
- Block Server Performance subprofile
- Copy Services subprofile
- Job Control subprofile
- Storage Enclosure subprofile (if expansion enclosures are attached)
- Disk Sparing subprofile
- Object Manager Adapter subprofile

- Thin Provisioning profile
- Pools from Volumes profile

The embedded SMI-S provider supports:

- HTTPS using SSL encryption on the default port 5989, or standard HTTP on the default port 5988. Both ports cannot be enabled at the same time.
- SLPv2
- CIM Alert and Lifecycle indications
- Microsoft Windows Server 2012 Server Manager and System Center Virtual Machine Manager

SMI-S implementation

SMI-S is implemented with the following components:

- CIM server (called a CIM Object Manager or CIMOM), which listens for WBEM requests (CIM operations over HTTP/HTTPS) from a CIM client, and responds.
- CIM provider, which communicates to a particular type of managed resource (for example, MSA 1050/2050 storage systems), and provides the CIMOM with information about them. In theory, providers for multiple types of devices (for example, MSA 1050/2050 storage systems and Brocade switches) can be plugged into the same CIMOM. However, in practice, all storage vendors provide the CIMOM and a single provider together, and they do not co-exist well with solutions from other vendors.

These components may be provided in several different ways:

- Embedded agent: The hardware device has an embedded SMI-S agent. No other installation of software is required to enable management of the device.
- SMI solution: The hardware or software ships with an agent that is installed on a host. The agent needs to connect to the device and obtain unique identifying information.

SMI-S architecture

The architecture requirements for the embedded SMI-S Array provider are to work within the Management Controller (MC) architecture, use limited disk space, use limited memory resources and be as fast as a proxy provider running on a server. The CIMOM used is the open source SFCB CIMOM.

SFCB is a lightweight CIM daemon that responds to CIM client requests and supports the standard CIM XML over [http/https](http://https) protocol. The provider is a CMPI (Common Management Protocol Interface) provider and uses this interface. To reduce the memory footprint, a third-party package called CIMPLE is used. For more information on SFCB go to <http://sourceforge.net/projects/sblim/files/sblim-sfcb>.

About the MSA 1050/2050 SMI-S provider

The provider is a SMI-S 1.5 provider which passes CTP 1.5 tests. Full provisioning is supported.

The MSA 1050/2050 SMI-S provider is a full-fledged embedded provider implemented in the firmware. It provides an industry-standard WBEM-based management framework. SMI-S clients can interact with this embedded provider directly and do not need an intermediate proxy provider. The provider supports active management features such as RAID provisioning.

The MSA 1050/2050 SAN and SAS system is supported. The classes for HPE are `HP_XXX`. The device namespace for HPE is `/root/hpq`.

The embedded CIMOM can be configured either to listen to secure SMI-S queries from the clients on port 5989 and require credentials to be provided for all queries, or to listen to unsecure SMI-S queries from the clients on port 5988. This provider implementation complies with the SNIA SMI-S specification version 1.5.0.

NOTE: Port 5989 and port 5988 cannot be enabled at the same time.

The namespace details are given below.

- Implementation Namespace - root/hpq
- Interop Namespace - root/interop

The embedded provider set includes the following providers:

- Instance Provider
- Association Provider
- Method Provider
- Indication Provider

The embedded provider supports the following CIM operations:

- getClass
- enumerateClasses
- enumerateClassNames
- getInstance
- enumerateInstances
- enumerateInstanceNames
- associators
- associatorNames
- references
- referenceNames
- invokeMethod

SMI-S profiles

SMI-S is organized around profiles, which describe objects relevant for a class of storage subsystem. SMI-S includes profiles for arrays, FC HBAs, FC switches, and tape libraries. Profiles are registered with the CIM server and advertised to clients using SLP. HPE SIM determines which profiles it intends to manage, and then uses the CIM model to discover the actual configurations and capabilities.

Table 20 Supported SMI-S profiles

Profile/subprofile/package	Description
Array profile	Describes RAID array systems. It provides a high-level overview of the array system.
Block Services package	Defines a standard expression of existing storage capacity, the assignment of capacity to Storage Pools, and allocation of capacity to be used by external devices or applications.
Physical Package package	Models information about a storage system's physical package and optionally about internal sub-packages.
Health package	Defines the general mechanisms used in expressing health in SMI-S.
Server profile	Defines the capabilities of a CIM object manager based on the communication mechanisms that it supports.
FC Initiator Ports profile	Models the Fibre Channel-specific aspects of a target storage system.
SAS Initiator Ports subprofile	Models the SAS-specific aspects of a target storage system.
iSCSI Initiator Ports subprofile	Models the iSCSI-specific aspects of a target storage system.
Access Points subprofile	Provides addresses of remote access points for management services.
Fan profile	Specializes the DMTF Fan profile by adding indications.
Power Supply profile	Specializes the DMTF Power Supply profile by adding indications.

Table 20 Supported SMI-S profiles (continued)

Profile/subprofile/package	Description
Profile Registration profile	Models the profiles registered in the object manager and associations between registration classes and domain classes implementing the profile.
Software subprofile	Models software or firmware installed on the system.
Masking and Mapping profile	Models device mapping and masking abilities for SCSI systems.
Disk Drive Lite profile	Models disk drive devices.
Extent Composition	Provides an abstraction of how it virtualizes exposable block storage elements from the underlying Primordial storage pool.
Location subprofile	Models the location details of product and its sub-components.
Sensors profile	Specializes the DMTF Sensors profile.
Software Inventory profile	Models installed and available software and firmware.
Storage Enclosure profile	Describes an enclosure that contains storage elements (e.g., disk or tape drives) and enclosure elements (e.g., fans and power supplies).
Multiple Computer System subprofile	Models multiple systems that cooperate to present a “virtual” computer system with additional capabilities or redundancy.
Copy Services subprofile	Provides the ability to create and delete local snapshots and local volume copies (clones), and to reset the synchronization state between a snapshot and its source volume.
Job Control subprofile	Provides the ability to monitor provisioning operations, such as creating volumes and snapshots, and mapping volumes to hosts.
Disk Sparing subprofile	Provides the ability to describe the current spare disk configuration, to allocate/de-allocate spare disks, and to clear the state of unavailable disk drives.
Object Manager Adapter subprofile	Allows the client to manage the Object Manager Adapters of a SMI Agent. In particular, it can be used to turn the indication service on and off.
Thin Provisioning profile	Specializes the Block Services Package to add support for thin provisioning of volumes. SMI-S does not support the creation of virtual pools. However, a client can create virtual volumes.
Pools from Volumes profile	Models a pool created from other volumes. This profile is used in conjunction with the Thin Provisioning profile to model virtual storage pools.

Block Server Performance subprofile

The implementation of the block server performance subprofile allows use of the CIM_BlockStorageStatisticalData classes and their associations, and the GetStatisticsCollection, CreateManifestCollection, AddOrModifyManifest and RemoveManifest methods.

The Block Server Performance subprofile collection of statistics updates at 60-second intervals.

CIM

Supported CIM operations

SFCB provides a full set of CIM operations including GetClass, ModifyClass, CreateClass, DeleteClass, EnumerateClasses, EnumerateClassNames, GetInstance, DeleteInstance, CreateInstance, ModifyInstance, EnumerateInstances, EnumerateInstanceNames, InvokeMethod (MethodCall), ExecQuery, Associators, AssociatorNames, References, ReferenceNames, GetQualifier, SetQualifier, DeleteQualifier, EnumerateQualifiers, GetProperty and SetProperty.

CIM Alerts

The implementation of alert indications allows a subscribing CIM client to receive events such as FC cable connects, Power Supply events, Fan events, Temperature Sensor events and Disk Drive events.

If the storage system's SMI-S interface is enabled, the system will send events as indications to SMI-S clients so that SMI-S clients can monitor system performance. For information about enabling the SMI-S interface, see [“SMI-S configuration” \(page 178\)](#).

In a dual-controller configuration, both controller A and B alert events are sent via controller A's SMI-S provider.

The event categories in [Table 21](#) pertain to FRU assemblies and certain FRU components.

Table 21 CIM Alert indication events

FRU/Event category	Corresponding SMI-S class	Operational status values that would trigger alert conditions
Controller	HP_Controller	Down, Not Installed, OK
Hard Disk Drive	HP_DiskDrive	Unknown, Missing, Error, Degraded, OK
Fan	HP_PSUFan	Error, Stopped, OK
Power Supply	HP_PSU	Unknown, Error, Other, Stressed, Degraded, OK
Temperature Sensor	HP_OverallTempSensor	Unknown, Error, Other, Non-Recoverable Error, Degraded, OK
Battery/Super Cap	HP_SuperCap	Unknown, Error, OK
FC Port	HP_FCPort	Stopped, OK
SAS Port	HP_SASTargetPort	Stopped, OK
iSCSI Port	HP_ISCSIEthernetPort	Stopped, OK

Life cycle indications

The SMI-S interface provides CIM life cycle indications for changes in the physical and logical devices in the storage system. The SMI-S provider supports all mandatory elements and certain optional elements in SNIA SMI-S specification version 1.5.0. CIM Query Language (CQL) and Windows Management Instrumentation Query Language (WQL) are both supported, with some limitations to the CQL indication filter. The provider supports additional life cycle indications that the Windows Server 2012 operating system requires.

Table 22 Life cycle indications

Profile or subprofile	Element description and name	WQL or CQL
Block Services	SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_StoragePool Send life cycle indication when a disk group is created or deleted.	Both
Block Services	SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_StorageVolume Send life cycle indication when a volume is created or deleted.	Both
Block Services	SELECT * FROM CIM_InstModification WHERE SourceInstance ISA CIM_LogicalDevice Send life cycle indication when disk drive (or any logical device) status changes.	Both
Copy Services	SELECT * FROM CIM_InstModification WHERE SourceInstance ISA CIM_StorageSynchronized AND SourceInstance.SyncState <> PreviousInstance.SyncState Send life cycle indication when the snapshot synchronization state changes.	CQL

Table 22 Life cycle indications (continued)

Profile or subprofile	Element description and name	WQL or CQL
Disk Drive Lite	SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_DiskDrive Send life cycle indication when a disk drive is inserted or removed.	Both
Job Control	SELECT * FROM CIM_InstModification WHERE SourceInstance ISA CIM_ConcreteJob AND SourceInstance.OperationalStatus=17 AND SourceInstance.OperationalStatus=2 Send life cycle indication when a create or delete operation completes for a volume, LUN, or snapshot.	WQL
Masking and Mapping	SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_AuthorizedSubject Send life cycle indication when a host privilege is created or deleted.	Both
Masking and Mapping	SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_ProtocolController Send life cycle indication when create/delete storage hardware ID (add/remove hosts).	Both
Masking and Mapping	SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_ProtocolControllerForUnit Send life cycle indication when a LUN is created, deleted, or modified.	Both
Multiple Computer System	SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_ComputerSystem Send life cycle indication when a controller is powered on or off.	Both
Multiple Computer System	SELECT * FROM CIM_InstModification WHERE SourceInstance ISA CIM_ComputerSystem AND SourceInstance.OperationalStatus <> PreviousInstance.OperationalStatus Send life cycle indication when a logical component degrades or upgrades the system.	WQL
Multiple Computer System	SELECT * FROM CIM_InstModification WHERE SourceInstance ISA CIM_RedundancySet AND SourceInstance.RedundancyStatus <> PreviousInstance.RedundancyStatus Send life cycle indication when the controller active-active configuration changes.	WQL
Target Ports	SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_FCPort Send life cycle indication when a target port is created or deleted.	Both
Target Ports	SELECT * FROM CIM_InstModification WHERE SourceInstance ISA CIM_FCPort AND SourceInstance.OperationalStatus <> PreviousInstance.OperationalStatus Send life cycle indication when the status of a target port changes.	WQL

SMI-S configuration

In the default SMI-S configuration:

- The secure SMI-S protocol is enabled, which is the recommended protocol for SMI-S.
- The SMI-S interface is enabled for the manage user.

Table 23 lists the CLI commands relevant to the SMI-S protocol:

Table 23 CLI commands for SMI-S protocol configuration

Action	CLI command
Enable secure SMI-S port 5989 (and disable port 5988)	<code>set protocols smis enabled</code>
Disable secure SMI-S port 5989	<code>set protocols smis disabled</code>
Enable unsecure SMI-S port 5988 (and disable port 5989)	<code>set protocols usmis disabled</code>
Enable unsecure SMI-S port 5988	<code>set protocol usmis enabled</code>
See the current status	<code>show protocols</code>
Reset all configurations	<code>reset smis-configurations</code>

To configure the SMI-S interface for other users:

1. Log in as manage
2. If the user does not already exist, create one using this command:
`create user role manage username`
3. Type this command:
`set user username interfaces wbi,cli,smis,ftp,sftp`

Listening for managed-logs notifications

For use with the storage system's managed logs feature, the SMI-S provider can be set up to listen for notifications that log files have filled to a point that are ready to be transferred to a log-collection system. For more information about the managed logs feature, see ["About managed logs" \(page 30\)](#).

To set up SMI-S to listen for managed logs notifications:

1. In the CLI, enter this command:
`set advanced-settings managed-logs enabled`
2. In an SMI-S client:
 - a. Subscribe using the `SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_LogicalFile` filter.
 - b. Subscribe using the `SELECT * FROM CIM_InstDeletion WHERE SourceInstance ISA CIM_LogicalFile` filter.

Testing SMI-S

Use an SMI-S certified client for SMI-S 1.5. HPE has clients such as HPE SIM and HPE Storage Essentials. Other common clients are Microsoft System Center, IBM Tivoli, EMC CommandCenter and CA Unicenter. Common WBEM CLI clients are Pegasus `cimcli` and Sblim's `wbemcli`.

To certify that the array provider is SMI-S 1.5 compliant, SNIA requires that the providers pass the Conformance Test Program (CTP) tests.

The `reset smis-configuration` command enables the restoration of your original SMI-S configuration.

Troubleshooting

Table 24 provides solutions to common SMI-S problems.

Table 24 Troubleshooting

Problem	Cause	Solution
Unable to connect to the embedded SMI-S Array provider.	SMI-S protocol is not enabled.	Log in to the array as manage and type: <code>set protocol smis enabled</code>
HTTP Error (Invalid username/password or 401 Unauthorized).	User preferences are configurable for each user on the storage system.	Check that the user has access to the <code>smis</code> interface and set the user preferences to support the <code>smis</code> interface, if necessary. See “Adding, modifying, and deleting users” (page 44) for instructions on how to add users. Also verify the supplied credentials.
Want to connect securely as user name <code>my_xxxx</code> .	Need to add user.	Log in to the array as manage. Type: <code>create user level manage my_xxxuser</code> and then type: <code>set user my_xxxuser interfaces wbi,cli,smis</code>
Unable to discover via SLP.	SLP multicast has limited range (known as hops).	Move the client closer to the array or set up a SLP DA server or using unicast requests.
Unable to determine if SMI-S is running.	Initial troubleshooting.	Install <code>wbemcli</code> on a Linux system by typing: <code>apt-get install wbemcli</code> Type: <code>wbemcli -nl -t -noverify ein 'https://manage:!manage@:5989/root/hpq:cim_computersystem'</code>
SMI-S is not responding to client requests.	SMI-S configuration may have become corrupted.	Use the CLI command <code>reset smis-configuration</code> . Refer to the CLI Reference Guide for further information.

Using SLP

MSA 1050/2050 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 25](#) and populate the configuration attributes shown in [Table 26](#).

Table 25 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
FTP/SFTP (firmware upgrade)	service:firmware-update:ftp/ sftp
SNMP	service:api:snmp

Table 26 SLP attributes shown for a storage system

SLP attribute	Corresponding property shown by the CLI <code>show system detail</code> 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

You can enable or disable the SLP service in the SMU, as described in [“Enabling or disabling system-management services” \(page 50\)](#), 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 that data for display and analysis. For information about the managed logs feature, see [“About managed logs” \(page 30\)](#).

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, SMI-S, 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 FTP/SFTP 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” \(page 164\)](#).


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.

Storing log files

It is recommended to store log files hierarchically by storage-system name, log-file type, and date/time. Then, if historical analysis is required, the appropriate log-file segments can easily be located and can be concatenated into a complete record.

For example, assume that the administrator of a log-collection system has created the following hierarchy for logs from two storage systems named Storage1 and Storage2:



In push mode, when the administrator receives an email with an attached ecdebug file from Storage1, the administrator would open the attachment and unzip it into the ecdebug subdirectory of the Storage1 directory.

In pull mode, when the administrator receives notification that an SC debug log needs to be transferred from Storage2, the administrator would use the storage system's FTP/SFTP interface to get the log and save it into the scdebug subdirectory of the Storage2 directory.

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. A storage system that uses only SSDs, without tiering.
all-flash array	See AFA.
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.
ALUA	Asymmetric Logical Unit Access.
array	See storage system.
ASC/ASCQ	Additional Sense Code/Additional Sense Code Qualifier. Information on sense data returned by a SCSI device.
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.
auto-write-through	See AWT.
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. See <i>also</i> compatible disk, dynamic spare, global spare.
AWT	Auto-write-through. A setting that specifies when the RAID controller cache mode automatically changes from write-back to write-through.
base volume	A virtual volume that is not a snapshot of any other volume, and is the root of a snapshot tree.
canister	See 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 sheetmetal housing of an enclosure.
child volume	The snapshot of a parent volume in a snapshot tree. See parent volume.
chunk size	The amount of contiguous data that is written to a disk group member before moving to the next member of the disk group.
CIM	Common Information Model. The data model for WBEM. It provides a common definition of management information for systems, networks, applications and services, and allows for vendor extensions.
CIMOM	Common Information Model Object Manager. A component in CIM that handles the interactions between management applications and providers.
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. See <i>also</i> available disk, dynamic spare, global spare.
controller A (or B)	A short way of referring to controller module A (or B).
controller enclosure	An enclosure that contains one or two controller modules.

controller module	A FRU that contains the following subsystems and devices: a Storage Controller processor; a Management Controller processor; a SAS expander and Expander Controller processor; management interfaces; cache protected by a supercapacitor pack and flash memory; host, expansion, network, and service ports; and midplane connectivity.
CPLD	Complex programmable logic device.
CQL	CIM Query Language.
CRC	Cyclic Redundancy Check.
CRU	customer-replaceable unit. 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. <i>See also</i> FRU.
CSV	Comma-separated values. A format to store tabular data in plain-text form.
DAS	Direct Attached Storage. A dedicated storage device that connects directly to a host without the use of a switch.
deallocation rate	The rate, in pages per minute, at which a virtual pool is deallocating pages from its volumes because they no longer need the space to store data.
default mapping	Host-access settings that apply to all initiators that are not explicitly mapped to that volume using different settings. <i>See also</i> explicit mapping, masking.
DES	Data Encryption Standard.
DHCP	Dynamic Host Configuration Protocol. A network configuration protocol for hosts on IP networks.
disk group	A group of disks that is configured to use a specific RAID level and provides storage capacity for a pool. <i>See also</i> virtual disk group, read cache.
drain	The automatic movement of active volume data from a virtual disk group to other disk-group members within the same pool.
drive enclosure	<i>See</i> expansion enclosure. <i>See also</i> EBOD, JBOD.
drive spin down	<i>See</i> DSD.
DSD	Drive spin down. A power-saving feature that monitors disk activity in the storage system and spins down inactive spinning disks based on user-selectable policies. Drive spin down is not applicable to disks in virtual pools.
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. <i>See also</i> available disk, compatible disk, global spare.
EBOD	Expanded Bunch of Disks. Expansion enclosure attached to a controller enclosure.
EC	Expander Controller. A processor (located in the SAS expander in each controller module and expansion module) that controls the SAS expander and provides SES functionality. <i>See also</i> EMP.
EEPROM	Electrically erasable programmable ROM.
EMP	Enclosure management processor. An Expander Controller subsystem that provides SES data such as temperature, power supply and fan status, and the presence or absence of disks.
enclosure	A physical storage device that contains I/O modules, disk drives, and other FRUs. <i>See also</i> controller enclosure, expansion enclosure.
enclosure management processor	<i>See</i> EMP.
ESD	Electrostatic discharge.
ESM	Environmental Service Module. <i>See</i> IOM.
Expander Controller	<i>See</i> EC.

expansion enclosure	An enclosure that contains one or two expansion modules. Expansion enclosures can be connected to a controller enclosure to provide additional storage capacity. <i>See also</i> 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.
explicit mapping	Access settings for an initiator to a volume that override the volume's default mapping. <i>See also</i> default mapping, masking.
failback	<i>See</i> 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. <i>See also</i> recovery.
FC	Fibre Channel.
FC-AL	Fibre Channel Arbitrated Loop. The FC topology in which devices are connected in a one-way loop.
FDE	Full Disk Encryption. A feature that secures all the user data on a storage system. <i>See also</i> lock key, passphrase, repurpose, SED.
FPGA	Field-programmable gate array. An integrated circuit designed to be configured after manufacturing.
FRU	field-replaceable unit. 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. Product interfaces use the term "FRU" to refer to both FRUs and CRUs. <i>See</i> CRU.
full disk encryption	<i>See</i> FDE.
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. <i>See also</i> available disk, compatible disk, dynamic spare.
HBA	Host bus adapter. A device that facilitates I/O processing and physical connectivity between a host and the storage system.
host	A user-defined group of initiators that represents a server.
host group	A user-defined group of hosts for ease of management, such as for mapping operations.
host port	A port on a controller module that interfaces to a host computer, either directly or through a network switch.
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.
I/O Manager	An SNMP MIB term for a controller module.
I/O module	<i>See</i> IOM.
IOM	Input/output module, or I/O module. An IOM can be either a controller module or an expansion module.
IOPS	I/O operations per second.
IQN	iSCSI Qualified Name.
iSCSI	Internet SCSI.
iSNS	Internet Storage Name Service.
JBOD	"Just a bunch of disks." <i>See</i> expansion enclosure.
LBA	Logical block address. The address used for specifying the location of a block of data.
leftover	The state of a disk that the system has excluded from a disk group because the timestamp in the 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 <i>also</i> FDE, passphrase.
loop	See FC-AL.
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	See MC.
map/mapping	Settings that specify whether a volume is presented as a storage device to a host system, and how the host system can access the volume. Mapping settings include an access type (read-write, read-only, or no access), controller host ports through which initiators may access the volume, and a LUN that identifies the volume to the host system. See <i>also</i> default mapping, explicit mapping, masking.
masking	A volume-mapping setting that specifies no access to that volume by hosts. See <i>also</i> default mapping, explicit mapping.
MC	Management Controller. A processor (located in a controller module) that is responsible for human-computer interfaces, such as the SMU, and computer-computer interfaces, such as SNMP, and interacts with the Storage Controller. See <i>also</i> EC, SC.
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. See <i>also</i> host, map/mapping, volume.
network port	The Ethernet port on a controller module through which its Management Controller is connected to the network.
NTP	Network time protocol.
NV device	Nonvolatile device. The CompactFlash memory card in a controller module.
OID	Object Identifier. In SNMP, an identifier for an object in a MIB.
orphan data	See unwritable cache data.
overcommit	A setting that controls whether a virtual pool is allowed to have volumes whose total size exceeds the physical capacity of the pool.
overcommitted	The amount of storage capacity that is allocated to virtual volumes exceeds the physical capacity of the storage system.
page	A range of contiguous LBAs in a virtual disk group.
paged storage	A method of mapping logical host requests to physical storage that maps the requests to virtualized “pages” of storage that are in turn mapped to physical storage. This provides more flexibility for expanding capacity and automatically moving data than the traditional, linear method in which requests are directly mapped to storage devices. Paged storage is also called virtual storage.
parent volume	A virtual volume that has snapshots (can be either a base volume or a base snapshot volume). The parent of a snapshot is its immediate ancestor in the snapshot tree.
partner firmware update	See PFU.
passphrase	A user-created password that allows users to manage lock keys in an FDE-capable system. See <i>also</i> 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 <i>also</i> 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. <i>See also</i> 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. <i>See also</i> peer connection.
PFU	Partner firmware update. The automatic update of the partner controller when the user updates firmware on one controller.
PGR	Persistent group reservations.
PHY	One of two hardware components that form a physical link between devices in a SAS network that enables transmission of data.
point-to-point	Fibre Channel Point-to-Point topology in which two ports are directly connected.
pool	<i>See</i> virtual pool.
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. <i>See also</i> replication set, secondary system.
primary volume	The volume that is the source of data in a replication set and that can be mapped to hosts. The primary volume exists in a primary pool in the primary storage system.
PSU	Power supply unit FRU.
quick rebuild	A virtual-storage feature that reduces the time that user data is less than fully fault-tolerant after a disk failure in a disk group. The quick-rebuild process rebuilds only data stripes that contain user data. Data stripes that have not been allocated to user data are rebuilt in the background.
RAID head	<i>See</i> controller enclosure.
RBOD	"RAID bunch of disks." <i>See</i> controller enclosure.
read cache	A special disk group, comprised of SSDs, that can be added to a virtual pool for the purpose of speeding up read access to data stored on spinning disks elsewhere in the pool. Read cache is also referred to as read flash cache.
read flash cache	<i>See</i> read cache.
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. <i>See also</i> failover.
remote syslog support	<i>See</i> syslog.
replication	Asynchronous replication of block-level data from a volume in a primary system to a volume in a secondary system by creating an internal snapshot of the primary volume and copying the snapshot data to the secondary system via Fibre Channel or iSCSI links. The capability to replicate volumes is a licensed feature.
replication set	For replication, a container that houses the infrastructure upon which replications are performed. It defines a relationship between a primary and secondary volume for the purposes of maintaining a remote copy of the primary volume on a peer system. <i>See</i> primary volume, secondary volume.
replication snapshot history	As part of handling a replication, the replication set will automatically take a snapshot of the primary and/or secondary volume, thereby creating a history of data that has been replicated over time. This feature can be enabled for a secondary volume or for a primary volume and its secondary volume, but not for a volume group.
repurpose	A method by which all data on a system or disk is erased in an FDE-capable system. Repurposing unsecures the system and disks without needing the correct passphrase. <i>See also</i> FDE, passphrase.
RFC	Read flash cache. <i>See</i> read cache.

SAS	Serial Attached SCSI.
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. <i>See also</i> EC, MC.
secondary system	The storage system that contains a replication set's secondary volume. <i>See also</i> replication set, primary system.
secondary volume	The volume that is the destination for data in a replication set and that is not accessible to hosts. The secondary volume exists in a secondary pool 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. <i>See also</i> FDE.
SEEPROM	Serial electrically erasable programmable ROM. A type of nonvolatile (persistent if power removed) computer memory used as FRU ID devices.
SES	SCSI Enclosure Services. The protocol that allows the initiator to communicate with the enclosure using SCSI commands.
SFCB	Small Footprint CIM Broker.
SFF	Small form factor.
SFTP	SSH File Transfer Protocol. A secure secondary interface for installing firmware updates, downloading logs, installing security certificates and keys, and installing a license. All data sent between the client and server will be encrypted.
SHA	Secure Hash Algorithm.
shelf	<i>See</i> enclosure.
SLP	Service Location Protocol. Enables computers and other devices to find services in a local area network without prior configuration.
SMART	Self-Monitoring Analysis and Reporting Technology. A monitoring system for disk drives that monitors reliability indicators for the purpose of anticipating disk failures and reporting those potential failures.
SMI-S	Storage Management Initiative - Specification. The SNIA standard that enables interoperable management of storage networks and storage devices. The interpretation of CIM for storage. It provides a consistent definition and structure of data, using object-oriented techniques.
SMU	Storage Management Utility. 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 both the source volume and in its pool. A snapshot can be mapped and written to. The capability to create snapshots is a licensed feature. A base of 64 snapshots is included without an additional license. Snapshots that can be mapped to hosts are counted against the snapshot-license limit, whereas transient and unmappable snapshots are not.
snapshot tree	A group of virtual volumes that are interrelated due to creation of snapshots. Since snapshots can be taken of existing snapshots, volume inter-relationships can be thought of as a "tree" of volumes. A tree can be 254 levels deep. <i>See also</i> base volume, child volume, parent volume, source volume.
SNIA	Storage Networking Industry Association. An association regarding storage networking technology and applications.
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	See SC.
Storage Management Utility	See SMU.
storage system	A controller enclosure with at least one connected expansion enclosure. Product documentation and interfaces use the terms storage system and system interchangeably.
syslog	A protocol for sending event messages across an IP network to a logging server. This feature supports User Datagram Protocol (UDP) but not Transmission Control Protocol (TCP).
thin provisioning	A virtual-storage feature that allows actual storage for a virtual volume to be assigned as data is written, rather than storage being assigned immediately for the eventual size of the volume. This allows the storage administrator to overcommit physical storage, which in turn allows the connected host system to operate as though it has more physical storage available than is actually allocated to it. When physical resources fill up, the storage administrator can add storage capacity on demand.
tier	<p>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:</p> <ul style="list-style-type: none"> • Performance, which uses SSDs (high speed) • Standard, which uses enterprise-class spinning SAS disks (10k/15k RPM, higher capacity) • Archive, which uses midline spinning SAS disks (<10k RPM, high capacity).
tier migration	The automatic movement of blocks of data, associated with a single virtual volume, between tiers based on the access patterns that are detected for the data on that volume.
tray	See enclosure.
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 for this term include unpresent and unmap.
unwritable cache data	Cache data that has not been written to disk and is associated with a volume that no longer exists or whose disks are not online. If the data is needed, the volume's disks must be brought online. If the data is not needed it can be cleared, in which case it will be lost and data will differ between the host system and disk. Unwritable cache data is also called orphan data.
UPS	Uninterruptible power supply.
UTC	Coordinated Universal Time.
UTF-8	UCS transformation format - 8-bit. A variable-width encoding that can represent every character in the Unicode character set used for the CLI andSMU interfaces.
virtual	The storage-class designation for logical components such as volumes that use paged-storage technology to virtualize data storage. See paged storage.
virtual disk group	A group of disks that is configured to use a specific RAID level. The number of disks that a virtual disk group can contain is determined by its RAID level. A virtual disk group can be added to a new or existing virtual pool. See <i>also</i> virtual pool.
virtual pool	A container for volumes that is composed of one or more virtual disk groups.
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 group	A user-defined group of volumes for ease of management, such as for mapping 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.
WBEM	Web-Based Enterprise Management.
WBI	Web-browser interface, called Storage Management Utility. The primary interface for managing the storage system. A user can enable the use of HTTP, HTTPS for increased security, or both. See SMU.
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

Symbols

* (asterisk) in option name 12

A

activity progress interface 69
all-flash array
 about 20
allocated space
 virtual storage 144
ALUA 27
archive tier 25
asterisk (*) in option name 12
Automated Tiered Storage
 about 25
 frequently accessed data 25
 infrequently accessed data 25

B

banner
 overview 139
base for size representations 15
bytes versus characters 15

C

cache
 configuring auto-write-through triggers and behaviors 76
 configuring host access to 75
 configuring system settings 75
 configuring volume settings 100
Capacity block
 physical and logical storage identification 39
capacity information 39
 viewing 143
Capacity Utilization panel 39, 144
certificate
 using FTP to install a security 171
 using SFTP to install a security 171
CHAP
 adding or modifying a record 84
 configuring for iSCSI hosts 84
 configuring through system settings 57, 58
 deleting a record 85
 overview 26
 setting up for use with a peer connection 124
 using in a system with a peer connection 84
 using with replication 124
characters versus bytes 15
color codes for storage space 14
compatibility
 when creating and modifying a peer connection 34
 when using replication queue policy 34
 when using replication snapshot history 34

configuration
 browser 11
 first-time 10
configuring system settings 41
contacting Hewlett Packard Enterprise 147
controllers
 restarting or shutting down 78
 using FTP to update firmware 166
 using SFTP to update firmware 166
 using the WBI to update firmware 66
Critical & Error Event Information panel 142
customer self repair 148

D

date and time
 configuring 140
 setting 41
debug data
 saving to a file 141
debug logs
 downloading 163
default mapping
 about 26
 advantages and disadvantages 111
DHCP
 configuring with system settings 49
disaster recovery
 accessing data from the backup system 120
 accessing data with intact replication set 120
 procedures 121
 using replication in 120
disk channels
 rescanning 64
disk groups
 about 17, 22
 adding 88
 configuring background scrub 77
 listed information 86
 modifying 90
 options 90
 read-cache 18, 89
 removing 91
 scrubbing 95
 virtual 18, 89
disk metadata
 clearing 65
disk settings
 configuring 73

- disks
 - configuring background scrub 77
 - configuring SMART 73
 - configuring spin down for available and global-spare 74
 - enabling reconstruction 28
 - scheduling spin down for all 74
 - using FTP to retrieve performance statistics 165
 - using FTP to update firmware 169
 - using SFTP to retrieve performance statistics 165
 - using SFTP to update firmware 169
 - using the WBI to update firmware 68
- documentation
 - providing feedback on 149
- drive spin down
 - configuring for available and global-spare disks 74
 - scheduling for all disks 74
- DWD (drive writes per day)
 - SSD endurance indicator 21
- dynamic spares
 - about 22
 - configuring 74

E

- EMP polling rate
 - configuring 73
- empty allocated pages
 - replication 119
- enclosure
 - front view 61
 - rear view 61
 - table view 62
 - viewing information about 61
- enclosure properties 62
- entering system identification information 52
- event log
 - viewing 143
- event notification
 - configuring 52
 - testing settings 55
- event severity icon 143
- explicit mapping
 - about 26, 111

F

- FDE
 - about 35
 - changing settings 70
 - clearing lock keys 71
 - repurposing disks 72
 - repurposing system 72
 - securing the system 71
 - setting FDE import lock key IDs 72
 - setting the passphrase 70

- firmware
 - about updating 30
 - updating through FTP 166
 - updating through SFTP 166
 - updating through the WBI 66
 - updating, best practices 66
 - using FTP to update controller module firmware 166
 - using FTP to update disk drive firmware 169
 - using FTP to update expansion module firmware 168
 - using SFTP to update controller module firmware 166
 - using SFTP to update disk drive firmware 169
 - using SFTP to update expansion module firmware 168
 - using the WBI to update controller module firmware 66
 - using the WBI to update disk firmware 68

- firmware update, monitoring progress 69

- firmware update, partner

- configuring 76

- footer

- overview 139

- foreign disk group

- resolving a resulting pool conflict 40

- FTP

- about updating firmware 166
 - downloading system heat map data 172
 - downloading system logs 163
 - retrieving disk-performance statistics 165
 - updating controller module firmware 166
 - updating disk drive firmware 169
 - updating expansion module firmware 168
 - using to install a security certificate 171
 - using with the log-management feature 164

- Full Disk Encryption

- See FDE

G

- global spares
 - adding and removing 92
- grouping
 - maximum number of hosts 26
 - maximum number of initiators 26
- guided setup
 - using 37

H

- heat map data
 - using FTP to download 172
 - using SFTP to download 172
- help
 - using online 13
- historical performance statistics
 - exporting 137
 - graphs 135
 - resetting 138
 - updating 137

- Home topic
 - host information 37
 - IOPS port information 38
 - port data throughput information 38
 - port information 38
 - spares information 40
 - storage capacity information 39
 - system health information 40
 - viewing system status 37
- host
 - adding initiators to 82
 - changing name 82
 - definition 38
 - removing initiators 82
 - viewing information about 80
- host access to cache
 - configuring 75
- host group
 - adding hosts 83
 - definition 38
 - removing hosts 83
 - renaming 83
- host groups
 - about 25
 - mapping 111
 - removing 83
 - viewing 80
- host I/O information
 - viewing 144
- host ports
 - changing settings 56
 - resetting 64
- hosts
 - about 25
 - adding to host group 83
 - basic information 80
 - list of 80
 - mapping 111
 - maximum number in a host group 26
 - removing 82
 - removing from host group 83

I

- icons
 - event severity 143
 - SMU communication status 140
- initiator
 - definition 37
 - deleting 81
 - manual creation 81
 - modifying 81
 - nickname 25

- initiators
 - about 25
 - adding to a host 82
 - mapping 111, 112
 - maximum number in a host 26
 - removing a mapping 114
 - removing all mappings 114
 - removing from a host 82
 - viewing 80
- iSCSI host security 26
- iSCSI IP version
 - configuring through System Settings 57, 59

J

- jumbo frames
 - configuring through System Settings 57, 59

L

- large pools
 - snapshot limits 104
- LDAP
 - about 31
 - user groups 46
- leftover disk 65
- licensed features
 - about 48
 - installing a license 48
 - installing a permanent license 49
 - snapshot limit 27
 - using FTP to install license file 171
 - using SFTP to install license file 171
 - viewing status of 49
 - volume copy 28
- lock key 35
- log data
 - saving to a file 141
- log management
 - about 30
 - using FTP 164
 - using SFTP 164
- log-collection system
 - administering 182
- logs
 - downloading debug 163
- LUNs
 - about 27
 - configuring response to missing 75

M

- managed logs
 - about 30
 - administering a log-collection system 182
 - enabling/disabling 78
 - pull mode 30
 - push mode 30

- mapping volumes
 - See volume mapping
- metadata
 - clearing disk 65
- MIB
 - See SNMP
- missing LUN response
 - configuring 75

N

- network ports
 - configuring with system settings 49
- nickname
 - initiator 25
- notification history
 - viewing 145
- notification settings
 - about 52
 - changing email 53
 - changing managed logs 54
 - sending SNMP 54
 - setting syslog 55
- NTP
 - configuring 140

O

- overcommitment setting
 - enabling 93
- overcommitting physical storage
 - about 24

P

- partner firmware update
 - configuring 76
- passphrase 35
- peer connections
 - CHAP setup 124
 - compatibility between MSA 1050/2050 and 1040/2040 systems 34
 - creating 123
 - deleting 126
 - modifying 125
 - querying 123
 - table 121
- performance monitoring
 - See storage system component performance
- performance statistics
 - about 29
 - historical performance graphs 135
 - resetting 138
 - viewing 135
- performance tier 25

- pools
 - about 22
 - attributes 86
 - changing settings 93
 - large pools snapshot limits 104
 - list of 86
 - viewing information about 86
 - virtual 22

- ports
 - attributes and status 38
 - data throughput 38
 - IOPS information 38

- priority
 - configuring utility 77

- provisioning
 - first-time 10

Q

- queue policy
 - compatibility between MSA 1050/2050 and 1040/2040 systems 34
 - setting from the Replications topic 127
 - setting from the Volumes topic 106
- quick rebuild
 - about 29

R

- read cache
 - about 21
 - advantages 21
 - cache utilization graph 39
- read-ahead caching
 - Adaptive option 24
 - Disabled option 24
 - optimizing 24
 - Stripe option 24
- read-cache disk groups
 - about 18
- reconstruction
 - about 28
- regulatory information 149
- remote support 148
- replication
 - about 116
 - between MSA 1050/2050 and 1040/2040 systems 34
 - creating a virtual pool for 119
 - of empty allocated pages 119
 - prerequisites 116
 - process 117
 - setting up snapshot space management for 119
 - using CHAP with 124
 - using in disaster recovery 120

- replication process 117
 - initial replication 117
 - internal snapshot space 119
 - subsequent replications 118
- replication sets
 - creating from the Replications topic 126
 - creating from the Volumes topic 105
 - deleting 130
 - maintaining snapshot history 127
 - modifying 129
 - primary volumes and volume groups 126
 - queuing from the Replications topic 127
 - queuing from the Volumes topic 106
 - secondary volumes and volume groups 127
- replication snapshot history
 - compatibility between MSA 1050/2050 and 1040/2040 systems 34
 - maintaining from the Replications topic 127
 - maintaining from the Volumes topic 106
- replications
 - initiating from the Replications topic 130
 - initiating from the Volumes topic 108
 - Peer Connections table 121
 - Replication Sets table 122
 - Replication snapshot history table 122
 - scheduling from the Replications topic 130
 - scheduling from the Volumes topic 108
 - viewing 121
- repurposing
 - disks 72
 - secured disks and systems 35
 - system 72
- rescan disk channels 64
- reserved space 144
- resetting host ports 64
- restarting controllers 78
- rolling back volume data
 - about 103
 - virtual volumes and snapshots 103

S

- SAS cabling
 - about 36
- SAS fan-out cables
 - configuration indicators 62
 - overview 36
 - usage indication 38
- schedules
 - deleting 60, 110, 134
 - managing 59
 - modifying from the Home topic 59
 - table information 98

- scheduling
 - replications from the Replications topic 130
 - replications from the Volumes topic 108
 - snapshot 104
 - snapshot reset 105
- scrub
 - configuring background disk 77
 - configuring background for disk groups 77
- SCSI MODE SELECT command
 - configuring handling of 75
- SCSI SYNCHRONIZE CACHE command
 - configuring handling of 75
- security certificate
 - using FTP to install 171
 - using SFTP to install 171
- SFTP
 - about updating firmware 166
 - downloading system heat map data 172
 - downloading system logs 163
 - retrieving disk-performance statistics 165
 - updating controller module firmware 166
 - updating disk drive firmware 169
 - updating expansion module firmware 168
 - using to install a security certificate 171
 - using with the log-management feature 164
- shutting down controllers 78
- sign out, auto
 - viewing remaining time 12
- signing in to the SMU 15
- single-controller system data-protection tips 35
- size representations
 - about 15
- SLP
 - advertised interfaces 181
 - attributes 181
 - overview 180
- SMART
 - configuring 73
- SMI-S
 - architecture 174
 - Array profile supported profiles and subprofiles 173
 - Block Server Performance subprofile 176
 - CIM alerts 177
 - components 173
 - configuring 178
 - embedded array provider 173
 - implementation 174
 - life cycle indications 177
 - managed-logs notifications 179
 - MSA 2040 SMI-S provider 174
 - profile descriptions 175
 - supported CIM operations 176
 - testing 179
 - troubleshooting 180

SMU

- about 10
- communication status icon 140
- signing in 15
- tips for using 12

snapshot

- basic information 96
- copying from the volumes topic 100
- creating 104
- resetting to current data in source volume 105

snapshot history

- maintaining from the Replications topic 127
- maintaining from the volumes topic 106

snapshot space management

- setting up in the context of replication 119

snapshots

- about 27
- about reset snapshot 28
- about rollback 28
- automatic deletion 27
- deleting 103
- list of 96
- list of child snapshots 96
- mapping 111
- resetting 27
- rollback feature 28
- setting snapshot pool space 27

snapshots, virtual

- levels 27
- parent-child relationships 27
- rollback feature 27
- snapshot hierarchy 27

SNMP

- configuring traps 159
- differences between FA MIB 2.2 and 4.0 162
- enterprise trap MIB 160
- enterprise traps 150
- external details for connUnitPortTable 159
- external details for connUnitRevsTable 156
- external details for connUnitSensorTable 158
- FA MIB 2.2 behavior 151
- FA MIB 2.2 objects, descriptions, and values 151
- management 159
- MIB-II behavior 150
- overview 150
- setting event notification 159

sorting a table 12

spares

- about 22
- Home topic information 40
- managing 92
- rules for 22
- types of 22

SSD read cache

- about 21

SSDs

- about 20
- cost/benefit analysis 20
- data retention 21
- DWD (drive writes per day) 21
- endurance indicated by DWD 21
- gauging percentage of life remaining 20
- internal disk management 21
- maintenance 20
- overprovisioning 21
- rules for using 20
- SSD Life Left disk property 20
- wear leveling 21

standard tier 25

storage blocks

- read cache 39
- virtual storage 39

Storage Management Utility

- about 10

storage system

- See system

storage system component performance

- about monitoring historical data 29

support

- Hewlett Packard Enterprise 147

synchronize-cache mode

- configuring 75

system

- data-protection tips for a single-controller 35

system activity

- viewing 144, 145

system components

- properties 62

system health 40

- viewing the health panel 141

system information

- configuring settings 52
- menu options 139
- viewing 139

System Information panel 139

system services

- configuring 50
- enabling or disabling 50

system settings

- changing host port settings 56
- changing system information settings 52
- changing system services 50
- configuring 41
- configuring CHAP 57, 58
- configuring network ports 49
- installing a license 48
- sending notifications 52
- sending SNMP notifications 54
- setting date and time 41

- system status
 - viewing 37
- system utilities
 - configuring 77

T

- table sorting 12
- tables
 - tips for using 12
- task schedule
 - See schedule
- technical support 147
- temperature
 - configuring controller shutdown for high 76
- thin provisioning
 - about 24
 - overcommit storage 24
- tiers
 - archive 25
 - performance 25
 - standard 25
 - viewing I/O information 144
- time and date
 - configuring 140
- troubleshooting resources 143

U

- ULP 27
- unallocated space
 - virtual storage 144
- units for size representations 15
- User Information panel 141
- user interface
 - main areas 11
- user panel
 - changing user settings 141
- users
 - adding 44
 - deleting 44
 - LDAP 46
 - local 42
 - modifying 44
- using guided setup 37
- utility priority
 - configuring 77

V

- virtual disk groups
 - about 18
 - number allowed per pool 18
 - removal requirements 91
 - requirements 18

- virtual pools
 - about 22
 - about removing 22
 - changing settings 93
 - viewing information about 86
 - volume allocation 22
- virtual snapshots
 - about 27
 - creation process 27
- virtual storage
 - about 16
 - advantages 16
 - page definition 16
 - quick rebuild 29
 - reconstruction using quick-rebuild features 28
- virtual volumes
 - about adding to virtual pools 22
 - creating 99
- volume cache options
 - about 23
- volume copy
 - aborting 101
 - about 28
- volume groups
 - about 23
 - adding volumes 101
 - mapping 111
 - maximum number of volumes 23
 - removing 102
 - removing group and volumes 102
 - removing volumes from 101
 - renaming 102
 - requirements 23
- volume mapping
 - about 26
 - defaults 26
 - editing 113
 - procedure 112
 - Related Maps table 80
 - unmapping 113
 - viewing details 114
 - viewing information about 80, 111
- volume tier affinity
 - about 25

volumes

- aborting a copy operation 101
- about 23
- about cache options 23
- adding to volume group 101
- basic information 96
- changing name 100
- configuring cache settings 100
- copying 100
- creating a virtual volume 99
- deleting 103
- expanding 100
- list of 96
- mapping 111
- modifying 100
- removing a mapping 114
- removing all mappings 114
- removing from a volume group 101
- rolling back data 103
- viewing information about 96
- virtual 23

Volumes topic

- Maps table 97
- Replication Sets table 97
- Schedules table 98
- Snapshots table 96
- Volumes table 96

W

- warranty information 148
 - HPE Enterprise servers 148
 - HPE Networking products 149
 - HPE ProLiant and x86 Servers and Options 148
 - HPE Storage products 149
- web-browser buttons to avoid 12
- web-browser setup 11
- websites
 - customer self repair 148
- welcome panel
 - using 37
- write-back caching 23
- write-through caching 23
- www.hpe.com/support/msa1050QuickSpecs 56
- www.hpe.com/support/msa2050QuickSpecs 56