



ThinkSystem DS6200/DS4200/DS2200/DS EXP Storage Manager Guide



Machine Type: 4619/4617/4599/4588

Firmware release: G250

Part Number: 01GW902

First Edition (May 2017)

© Copyright Lenovo 2017.

LIMITED AND RESTRICTED RIGHTS NOTICE: If data or software is delivered pursuant a General Services Administration "GSA" contract, use, reproduction, or disclosure is subject to restrictions set forth in Contract No. GS-35F-05925.

Lenovo, the Lenovo logo, BladeCenter, DS Series, Flex System, NeXtScale System, System x, and ThinkSystem are trademarks of Lenovo in the United States, other countries, or both. Other company, product, or service names may be trademarks or service marks of others.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Contents

About this guide.	10
Intended audience.	10
Prerequisites	10
Related documentation	10
Document conventions and symbols	11
1 Getting started	12
Configuring and provisioning a new storage system.	12
Using the interface	12
Web browser requirements and setup.	12
Areas of the interface	12
Tips for using the SMC	13
Tips for using tables	14
Tips for using help	14
Color codes	15
Size representations	16
Signing in and signing out	17
System concepts	17
About virtual storage	17
About disk groups	18
About RAID levels	19
About SSDs	20
About SSD read cache	22
About spares	22
About pools	22
About volumes and volume groups	23
About volume cache options	23
About thin provisioning	25
About automated tiered storage	25
About initiators, hosts, and host groups	26
About volume mapping	27
About snapshots	28
About copying volumes	29
About reconstruction	29
About quick rebuild	29
About performance statistics	30
About firmware update	30
About managed logs	31
About replicating virtual volumes	32
About the Full Disk Encryption feature	32
About data protection with a single controller	33
2 Working in the Home topic.	34
Viewing overall system status	34
Host information	34
Port information	34
Capacity information	35
Storage information	35

System health information	36
Spares information	36
Resolving a pool conflict caused by inserting a foreign disk group	36
Using the Configuration Wizard	36
Using the Configuration Wizard: Introduction	37
Using the Configuration Wizard: Set Date and Time	37
Using the Configuration Wizard: Password Setup	38
Using the Configuration Wizard: Network configuration	38
Using the Configuration Wizard: Enable system-management services	39
Using the Configuration Wizard: System information	40
Using the Configuration Wizard: Configure event notification	40
Using the Configuration Wizard: Port configuration	41
Using the Configuration Wizard: Confirm the configuration changes	43
Changing system information settings	43
Managing users	43
User options	44
Adding, modifying, and deleting users	45
Changing notification settings	46
Managing scheduled tasks	48
Installing a license	49
Viewing the status of licensed features	49
Installing a permanent license	49
3 Working in the System topic	50
Viewing system components	50
Front view	50
Rear view	50
Table view	51
Managing global spares	52
Changing system services settings	53
Changing network interface settings	54
Changing host port settings	55
Resetting host ports	57
Rescanning disk channels	57
Clearing disk metadata	57
Updating firmware	58
Best practices for firmware update	58
Updating controller module firmware	59
Updating expansion module firmware	60
Updating disk-drive firmware	60
Using the activity progress interface	61
Changing FDE settings	62
Changing FDE general configuration	62
Repurposing the system	64
Repurposing disks	64
Setting import lock key IDs	65
Restarting or shutting down controllers	65
Restarting controllers	65
Shutting down controllers	66

4 Working in the Hosts topic.....	67
Viewing hosts	67
Hosts table	67
Related Maps table.....	67
Creating an initiator	68
Modifying an initiator	68
Deleting initiators	68
Adding initiators to a host	68
Removing initiators from hosts.....	69
Removing hosts	69
Renaming a host	69
Adding hosts to a host group	69
Removing hosts from a host group	70
Renaming a host group	70
Removing host groups.....	70
Configuring CHAP.....	70
5 Working in the Pools topic	72
Viewing pools	72
Pools table	72
Related Disk Groups table.....	72
Related Disks table.....	73
Adding a disk group	74
Add Disk Group panel overview.....	74
Using SSDs in a disk group	75
Virtual disk groups.....	75
Read-cache disk groups	75
Disk group options.....	75
Modifying a disk group	76
Removing disk groups	77
Creating a volume	77
Changing pool settings	78
Verifying and scrubbing disk groups	79
Verifying a disk group.....	79
Scrubbing a disk group.....	79
6 Working in the Volumes topic	81
Viewing volumes	81
Volumes table	81
Snapshots table.....	81
Maps table.....	82
Replication Sets table	82
Schedules table.....	83
Creating a virtual volume	84
Modifying a volume.....	85
Adding volumes to a volume group.....	85
Removing volumes from a volume group	86
Renaming a volume group	86
Removing volume groups.....	86
Rolling back a volume.....	87
Deleting volumes and snapshots	88

Creating snapshots	88
Resetting a snapshot.....	89
Creating a replication set from the Volumes topic	90
Primary volumes and volume groups.....	90
Secondary volumes and volume groups	90
Initiating replication from the Volumes topic	91
7 Working in the Mappings topic.....	92
Viewing mappings	92
Mapping initiators and volumes	92
Viewing map details	95
8 Working in the Replications topic	96
About replicating virtual volumes.....	96
Replication prerequisites	96
Replication process.....	96
Creating a virtual pool for replication.....	99
Setting up snapshot space management in the context of replication	99
Replication and empty allocated pages.....	99
Disaster recovery.....	99
Replication licensing	101
Viewing replications	101
Peer Connections table	101
Replication Sets table.....	101
Creating a peer connection.....	102
CHAP and replication.....	103
Modifying a peer connection	104
Deleting a peer connection	104
Creating a replication set from the Replications topic.....	105
Primary volumes and volume groups.....	105
Secondary volumes and volume groups	105
Modifying a replication set	106
Deleting a replication set	106
Initiating replication.....	106
Scheduling replications.....	107
Aborting a replication	108
Suspending a replication	108
Resuming a replication.....	109
9 Working in the Performance topic.....	110
Viewing performance statistics	110
Historical performance graphs	110
Updating historical statistics	112
Exporting historical performance statistics.....	112
Resetting performance statistics.....	113
10 Working in the banner and footer	114
Banner and footer overview	114
Viewing system information	114
Viewing certificate information	114
Viewing connection information	115

Viewing system date and time information	115
Changing date and time settings.....	115
Viewing user information	116
Viewing health information	116
Saving log data to a file	116
Viewing event information.....	117
Viewing the event log	118
Viewing capacity information	118
Viewing host I/O information.....	119
Viewing tier I/O information.....	119
Viewing recent system activity.....	119
Viewing the notification history.....	120
A Other management interfaces.....	121
SNMP reference	121
Supported SNMP versions.....	121
Standard MIB-II behavior	121
Enterprise traps.....	121
FA MIB 2.2 SNMP behavior	122
External details for certain FA MIB 2.2 objects	127
Configuring SNMP event notification in the SMC	130
SNMP management.....	130
Enterprise trap MIB	130
Using FTP.....	132
Downloading system logs	132
Transferring log data to a log-collection system	133
Downloading historical disk-performance statistics	134
Updating firmware	135
Installing a license file	138
Installing a security certificate	139
Downloading system heat map data	139
Using SLP.....	140
B Administering a log-collection system	142
How log files are transferred and identified	142
Log-file details	142
Storing log files.....	142
C Best practices	144
Pool setup	144
RAID selection.....	144
Disk count per RAID level	144
Disk groups in a pool	145
Tier setup.....	145
VMware missing LUN response	145
Multipath configuration.....	145
Physical port selection.....	146
Boot from SAN	147

Glossary..... 148

Index155

Tables

1	Related documentation	10
2	Document conventions	11
3	Areas of the SMC interface.....	13
4	Home topic storage space color codes.....	15
5	Create Virtual Volumes panel storage space color codes.....	16
6	Storage size representations in base 2 and base 10.....	16
7	Decimal (radix) point character by locale	17
8	Example applications and RAID levels	19
9	RAID level comparison.....	20
10	Number of disks per RAID level to optimize virtual disk group performance	20
11	Settings for the default users.....	44
12	Additional information for rear view of enclosure.....	50
13	Activity progress properties and values.....	61
14	Available host groups, hosts, and initiators.....	93
15	Available volume groups and volumes.....	94
16	Historical performance graphs	110
17	Connection information	115
18	FA MIB 2.2 objects, descriptions, and values	122
19	connUnitRevsTable index and description values	127
20	connUnitSensorTable index, name, type, and characteristic values	128
21	connUnitPortTable index and name values	129
22	Interfaces advertised by SLP	140
23	SLP attributes shown for a storage system	140
24	RAID level characteristics and use cases	144
25	Recommended disk group sizes	145

About this guide

This guide provides information about managing a Lenovo ThinkSystem™ DS6200/DS4200/DS2200/DS EXP storage system by using its web interface, Storage Management Console (SMC).

Intended audience

This guide is intended for storage system administrators.

Prerequisites

Prerequisites for using this product include knowledge of:

- Network administration
- Storage system configuration
- Storage area network (SAN) management and server-attached storage
- Fibre Channel (FC) protocol
- Serial Attached SCSI (SAS) protocol
- Internet SCSI (iSCSI) protocol
- Ethernet protocol

Related documentation

Table 1 Related documentation

For information about	See
Obtaining printed documentation	Lenovo customer letter*
Obtaining multi-language safety information, environmental notices, warranties, service and support, licenses, and product documentation	Lenovo Read Me First*
Overview of hardware installation	Lenovo Storage DS6200/DS4200/DS2200/DS EXP Getting Started*
Product hardware installation and maintenance	Lenovo Storage DS6200/DS4200/DS2200/DS EXP Hardware Installation and Maintenance Guide
Using the command-line interface (CLI) to configure and manage the product	Lenovo Storage DS6200/DS4200/DS2200/DS EXP CLI Reference Guide
Event codes and recommended actions	Lenovo Storage DS6200/DS4200/DS2200/DS EXP Event Descriptions Reference Guide
Enhancements, known issues, and late-breaking information not included in product documentation	Lenovo Storage Firmware Release Notes

* Printed document included in product shipkit.


To obtain PDF versions of product documentation, visit <http://support.lenovo.com>.

Document conventions and symbols

Table 2 Document conventions

Convention	Element
Colored text	Cross-reference links
Black, underlined text	Email addresses
Colored, underlined text	Website addresses
Bold text	<ul style="list-style-type: none">Keys that are pressedText typed into a GUI element, such as a boxGUI elements that are clicked or selected, such as menu and list items, buttons, and check boxes
<i>Italic</i> text	Text emphasis
Monospace text	<ul style="list-style-type: none">File and directory namesSystem outputCodeCommands, their arguments, and argument values
<i>Monospace, italic</i> text	<ul style="list-style-type: none">Code variablesCommand variables
Monospace, bold text	Emphasis of file and directory names, system output, code, and text typed at the command line

 **CAUTION:** Indicates that failure to follow directions could result in damage to equipment or data.

 **IMPORTANT:** Provides clarifying information or specific instructions.

NOTE: Provides additional information.

 **TIP:** Provides helpful hints and shortcuts.

1 Getting started

The Storage Management Console (SMC) is a web-based application for configuring, monitoring, and managing the storage system. It 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 SMC. In a dual-controller system, 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 SMC, each controller module in the storage system has SNMP, FTP, 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.

Configuring and provisioning a new storage system

To configure and provision a storage system for the first time

1. Configure your web browser to use the SMC as described in [“Web browser requirements and setup” \(page 12\)](#).
2. Sign in to the SMC; the default user for management is `manage` and the default password is `!manage`. For more information about signing in, see [“Signing in and signing out” \(page 17\)](#).
3. Verify that controller modules and expansion modules have the latest firmware as described in [“Updating firmware” \(page 58\)](#).
4. Use the Configuration Wizard as described in [“Using the Configuration Wizard” \(page 36\)](#).
5. Create disk groups and pools, as described in [“Adding a disk group” \(page 74\)](#).
6. Create volumes and map them to initiators, as described in [“Creating a volume” \(page 77\)](#).
7. From hosts, verify volume mappings by mounting the volumes and performing read/write tests to the volumes.
8. Optionally, for replication of virtual volumes and snapshots, create peer connections and replication sets, as described in [“Creating a peer connection” \(page 102\)](#), [“Creating a replication set from the Replications topic” \(page 105\)](#), and [“Creating a replication set from the Volumes topic” \(page 90\)](#).

Using the interface

Web browser requirements and setup

- Use Mozilla Firefox 11 and newer, Google Chrome 17 and newer, Microsoft Internet Explorer 10 and 11, or Apple Safari 5.1 and newer.
- 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 SMC is configured to use HTTPS, ensure that Internet Explorer is set to use either TLS 1.0, TLS 1.1 or TLS 1.2.

Areas of the interface



The main areas of the SMC 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 3 Areas of the SMC interface

Banner:	Product ID	System panel (page 114)	Connection panel (page 115)	Date/time panel (page 115)	User panel (page 116)	Sign Out button (page 17)	Help button (page 14)
Topic tabs:	Home (page 34)	<div>Topic pane</div>					
	System (page 50)						
	Hosts (page 67)						
	Pools (page 72)						
	Volumes (page 81)						
	Mapping (page 92)						
	Replications (page 96)						
	Performance (page 110)						
Footer:	Health panel (page 116)	Event panel (page 117)	Capacity panel (page 118)	Host I/O panel (page 119)	Tier I/O panel (page 119)	Activity panel (page 119)	

Tips for using the SMC

- Do not use the browser's Back, Forward, Reload, or Refresh buttons. The SMC has a single page for which content changes as you perform tasks and automatically updates to show current data.
- A red asterisk (*) identifies a required setting.
- As you set options in action panels, the SMC 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 SMC and the controller you are accessing goes offline, the system informs you that the system is unavailable or that communication has been lost. After the controller comes back online, close and reopen the browser and start a new SMC session.
- If your session is inactive for too long, you will be signed out automatically. This timer resets after each action you perform. One minute before automatic sign-out you will be prompted to continue using the SMC.
- 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.

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.

Tips for using help

- To display help for the content in the topic pane, click the help icon (?) in the banner.
- 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 4 Home topic storage space color codes


















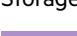





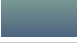
Color	Meaning
System performance statistics	
	IOPS
	Data throughput (MB/s)
Capacity graph, bottom bar	
	System physical space available
	System physical space used by global spares
	System physical space used by virtual disk groups
Capacity graph, top bar	
	Virtual pool reserved space (RAID parity and 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

Table 4 Home topic storage space color codes (continued)

Color	Meaning
Spares	
	Standard tier global spares
	Archive tier global spares

Create Virtual Volumes panel

Table 5 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, Portuguese, and Spanish
- 2 bytes per character for Arabic and Russian
- 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 SMC, 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.

Table 6 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 7 Decimal (radix) point character by locale

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

NOTE: Arabic, Portuguese, and Russian are not supported in this release.

Signing in and signing out

Multiple users can be signed in to each controller simultaneously.

For each active SMC session, an identifier is stored in the browser. 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 SMC session, but all instances of Firefox, Chrome, and Safari share the same SMC 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 10.14.33 and not 10.14.033.) The SMC Sign In page is displayed. If the Sign In page does not display, verify that you have entered the correct IP address.
2. On the sign-in page, enter the name and password of a configured user. The default user name and password are `manage` and `!manage`. To display the interface in a language other than the user setting, select the language from the Language list.

Language preferences can be configured for the system and for individual users.

3. Click **Sign In**. If the system is available, the Home page is displayed. 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 SMC 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 25](#), automated tiered storage on [page 25](#), read cache on [page 22](#), and the quick rebuild feature on [page 29](#).

About disk groups


A *disk group* is an aggregation of disks of the same type, using a specific RAID type 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).

 **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.

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

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 a combination of SSDs and spinning disks can only be created with 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 within a virtual group should have the same RAID level, disk capacity, 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. Otherwise, the data will be lost. When the last disk group is removed, the pool ceases to exist, and will be deleted from the system automatically.

The RAID type for a virtual disk group must be fault tolerant. The supported RAID types 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 a single disk in an NRAID configuration. For more information on read cache, see [“About SSD read cache” \(page 22\)](#).

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 8 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 9 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 10 Number of disks per RAID level to optimize virtual disk group performance

RAID level	Number of disks (data and parity)
1	N/A. 2 total.
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	16 total (8 RAID-1 subgroups)

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. You can create disk groups comprised of only SSDs, only spinning disks, or a combination of both. When combined with virtual disk groups that consist of other classes of disks, improved read and write performance is possible through automated tiered storage. You can also use one or two SSDs in read-cache disk groups to increase performance for pools without a Performance Tier license. For more information about automated tiered storage, see [“About automated tiered storage” \(page 25\)](#). For more information on read-cache disk groups, see [“Read-cache disk groups” \(page 19\)](#).

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” \(page 21\)](#).

- 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.
- 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 both SSDs and spinning disks.

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

The all-flash array feature, enabled by default, allows systems to run exclusively with disk groups that consist of SSDs, providing the ability to have a homogeneous SSD-only configuration. Systems using all-flash array have one tier that consists solely of SSDs. If a system includes disk groups with spinning disks, the disk groups must be removed before the all-flash array feature 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, over-provisioning to minimize write amplification, and support for Unmap commands.

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.

TRIM and UNMAP commands

A command (known as TRIM in the ATA command set and UNMAP in the SCSI command set) allows an operating system to inform an SSD of the blocks of data that are no longer considered in use and can be wiped internally. The autonomic nature of Lenovo ThinkSystem real-time tiering does not support static use of the TRIM or UNMAP commands. Lenovo ThinkSystem firmware does not issue TRIM or UNMAP commands to SSDs.

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 DWPDP 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 Lenovo ThinkSystem real-time tiering, can yield dramatic performance improvements.

Since Lenovo ThinkSystem 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, as well as the use of higher DWPDP SSDs (such as 8 DWPDP or 10 DWPDP). For less demanding application environments, the cost savings of upcoming 3 DWPDP SSDs may be more attractive.

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

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 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.
- SSDs do 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 RAID. 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 the system. A controller automatically reconstructs a fault-tolerant disk group (RAID 1, 5, 6, 10) when one or more of its disks fails and a compatible spare disk is available. A disk is compatible if it has enough capacity to replace the failed disk and is the same speed and type (enterprise SAS, for example). If the disks in the system are FDE-capable, 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 available, 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 drives in the form 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 type that is incorporated as a component of a pool, that stores volume data. For virtual pools, which can have multiple disk groups, volumes are added to a pool and the data is distributed across the pool's disk groups.

If the owning controller fails, the partner controller assumes temporary ownership of disk groups and resources owned by the failed controller. If a fault-tolerant cabling configuration and 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” \(page 74\)](#).

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 300 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 300 TiB (1 PiB if overcommit is enabled), 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 512 TiB 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.

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 which to move the data. 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 the storage for a file system partition you create with your operating system or third-party tools. For more information about mapping, see [“About volume mapping” \(page 27\)](#).

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.


Volume groups

For ease of management related to virtual storage, you can group 1–20 virtual 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.

About volume cache options

You can set options that optimize reads and writes performed for each volume.

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.

You can change the write-back cache setting for a volume when modifying it. *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 operation and 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.

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.
- The **Stripe** option sets the read-ahead size to one stripe. The controllers treat RAID 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 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 (ATS) 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:

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

Some advantages of having multiple tiers that use ATS are:

- Because a virtual pool can have multiple disk groups, each belonging to a different tier, a virtual pool can provide multiple tiers of storage.
- The I/O load is automatically balanced between components in a tier.
- Virtual disk groups can be added or removed without disrupting I/O. Data in virtual disk groups that are being removed is automatically migrated to other disk groups as long as the other disk groups have enough storage space for it. If they do not have the space, the system will not delete the disk groups until enough data is removed.

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. 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.

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.
- 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, or a port in a network switch.

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. A maximum of 512 names can be assigned.

For ease of management, you can group 1–128 initiators that represent a server or switch 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 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.

Default mapping is expected by some operating systems, such as Microsoft Windows, which can immediately discover the volume. 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.


If multiple hosts mount a volume without being cooperatively managed, volume data is at risk for corruption. To control access by specific initiators, 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 an initiator 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 can only be 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 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. Therefore, it is recommended to use the same LUN for explicit mappings as for the default mapping.

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

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.

With a license, you can create additional snapshots. When you reach the maximum number of base 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. For the maximum number of snapshots for your system, see the System configuration limits topic of the Storage Management Console online help for your system.

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.

Virtual 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. This feature operates differently depending on the storage technology for the snapshot.

Virtual snapshots also use the reset snapshot feature, which enables you to replace the data in a snapshot with the current data in the source volume. You can use it to update an existing snapshot with the data contained in the current 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.

Automatically deleting snapshots is currently unavailable.

Virtual snapshots

The process of 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 snapshot is modified.) There are no intermediate steps needed like designating the volume for snapshot capability. 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, they 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 virtual 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 creates a complete “physical” copy of a source 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 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 using the CLI to create a copy of a virtual base volume or snapshot, see the CLI Reference Guide.

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 component. Component 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 Hardware Installation and Maintenance Guide.

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

When reconstruction is complete, 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 feature for virtual storage 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 from host to disk, 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 5 minutes 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” \(page 110\)](#), [“Updating historical statistics” \(page 112\)](#), [“Exporting historical performance statistics” \(page 112\)](#), and [“Resetting performance statistics” \(page 113\)](#).

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 at a customer maintenance depot 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” \(page 58\)](#). 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 before any data is lost. The transfer does not remove any data from the logs in the storage system. This feature is disabled by default.

The *log-collection system* is a host computer that is designated to receive the log data transferred from the storage system. Because log data is transferred incrementally, the log-collection system is responsible for integrating the log data for display and analysis.

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 or SNMP to the log-collection system, which can then use FTP 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 SMC or the `get_logs` command in the FTP 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” \(page 116\)](#). For information about using the FTP interface, see [“Using FTP” \(page 132\)](#).

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” \(page 96\)](#).

About the Full Disk Encryption feature

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 without needing the correct passphrase. Repurposing erases all data and unsecures the system and disks.

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” \(page 62\)](#).

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” \(page 57\)](#).

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, data is held in controller cache until it is written to disk. In write-through mode, data is written directly 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. This can result in data becoming unavailable or, in some cases, volume unavailability.

If the controller can be brought back online long enough to perform a proper shutdown, 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, you can move its CompactFlash cache card to a replacement controller. If the CompactFlash cache card contains the cache data, this enables the data to be available when the new controller comes online. The CompactFlash card is externally accessible from the back of the controller, as described in the User Guide for your product.

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 significant 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” \(page 85\)](#).

2 Working in 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 or switch. 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	FC port: Name, type, ID (WWN), status, configured speed, actual speed, topology, primary loop ID, supported speeds, SFP status, part number, and health 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 SAS port: Name, type, ID (WWN), status, actual speed, topology, expected lanes, active lanes, disabled lanes, and health
------------------	---

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 SMC 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.

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” \(page 15\)](#).

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 different meanings for virtual storage technologies.

- 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, 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. For virtual storage, the Storage A block shows information for pool A, which is owned by controller A. The Storage B block shows the same types of information about virtual pool B. In a single-controller system, only the storage block relevant to that controller will be shown (for example, only the Storage A block will be shown if controller A is the sole operating controller).

Each storage block contains color-coded graphs. For color-code descriptions, see [“Color codes” \(page 15\)](#).

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(s) 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 pool A on `OldSystem`, the imported pool A from `OldSystem` will be offline.

To avoid this, do either of the following:

- Physically remove all disks for the existing pool, which will remove the pool, and then insert the imported disks.

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

- Delete the existing pool and then insert the imported disks.

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

Either method will allow the system to create pool A for the new disk group without conflict, allowing the imported disk group's data to be accessible. 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.


Using the Configuration Wizard

The Configuration Wizard helps you initially configure the system or change system configuration settings. When you complete this wizard you are given the option to start creating disk groups.

Using the Configuration Wizard: Introduction

You can use the Configuration Wizard to perform the following:

- Change the system date and time settings
- Change passwords for the default users, providing they still exist
- Configure each controller's network port
- Enable or disable system-management services
- Enter information to identify the system
- Configure event notification
- Configure controller host ports (if applicable)

The wizard guides you through each step. As you complete a step, it is highlighted at the bottom of the panel. For each step, you can view help by clicking the help icon . At any point, you can cancel the wizard and discard changes.

To use the Configuration Wizard, perform one of the following:

- Point to the Home tab, and select **Configuration Wizard**.
- In the Home topic, select **Action > Configuration Wizard**.

When the Configuration Wizard panel opens, click **Next** to proceed to the next step.

Using the Configuration Wizard: Set Date and Time

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, such as for Daylight Saving Time. You must make such adjustments manually.

NOTE: If you make changes in this step, they will be applied when you click **Next**. Changes made in other steps will be applied when you complete the wizard.

To use manual date and time settings

1. Clear the **Network Time Protocol (NTP)** check box.
2. To set the Date value, enter the current date in the format **YYYY-MM-DD**.
3. To set the Time value, enter two-digit values for the hour and minutes and select either **AM**, **PM**, or **24H** (24-hour clock).
4. Click **Next** to proceed to the next step.

To obtain the date and time from an NTP server

1. Select the **Network Time Protocol (NTP)** check box.

2. Perform one of the following:
 - o To have the system retrieve time values from a specific NTP server, enter its 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.
3. 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.
4. Click **Next** to proceed to the next step.

Using the Configuration Wizard: Password Setup

The system provides the default users `manage` and `monitor`.

1. To secure the storage system, enter and confirm a new password for each default user.

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. It can include printable UTF-8 characters except for the following: a space or " , < > \
2. Click **Next** to proceed to the next step.

Using the Configuration Wizard: Network configuration

You can change addressing parameters for the network port in each controller module. You can set static IP values or use DHCP (enabled by default). When setting static IP values, use IPv4 format.

In 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, current addressing is unchanged. You must have some means of determining what addresses have been assigned, such as the list of bindings on the DHCP server.

Each controller has the following factory-default IP settings:

- IP address source: DHCP
- 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

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: 0.0.0.0

⚠ 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 use DHCP

1. Set IP address source to DHCP. The new IP values will not appear until the Configuration Wizard is completed and you have logged in again to the new IP addresses.
2. Record the new addresses.
3. Click **Next** to proceed to the next step.

To use static IP values

1. Determine the IP address, subnet mask, and gateway values to use for each network port.
2. Set IP address source to manual.
3. Enter IP address, subnet mask, and gateway values for each controller. You must set a unique IP address for each controller.

NOTE: The following IP addresses are reserved for internal use by the storage system: 192.168.200.253, 192.168.200.254, 172.22.255.253, 172.22.255.254, and 127.0.0.1

4. Record the IP values you assign.
5. Click **Next** to proceed to the next step.

Using the Configuration Wizard: Enable 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 SMC (the web browser interface), CLI, or FTP, see [“Adding, modifying, and deleting users” \(page 45\)](#).

To change system services settings

1. 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. You can enable use of HTTP and/or HTTPS for increased security. Also, if you choose to disable the SMC, the change does not take effect until the Configuration Wizard has finished and you have logged in again. If you disable both, you will lose access to this interface.
 - 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.
 - 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.
 - o File Transfer Protocol (FTP). A secondary interface for installing firmware updates, downloading logs, and installing a license.
 - 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.

NOTE: Properly shut down the debug console by using the CLI command `set protocols debug disable`. Do not just close the console directly or by using the CLI command `exit`.

- 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.
2. Click **Next** to proceed to the next step.

Using the Configuration Wizard: System information

To change system information settings

1. 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: " < > \
2. Click **Next** to proceed to the next step.

Using the Configuration Wizard: Configure event notification

You can enable the system to:

- Send notifications to SNMP trap hosts.
- Send notifications to email addresses when events occur in the system.
- Use the managed logs feature, which transfers log data to a log-collection system. For more information about the managed logs feature, see [“About managed logs” \(page 31\)](#).
- Set remote syslog notification options 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.

To change SNMP notification settings

1. Select one of the following Notification Level options:
 - o **none** (disabled). All events are excluded from trap notification and traps are disabled.
However, Critical events and managed-logs events are sent regardless of the notification setting.
 - o **Critical**. Notifications are sent for Critical events only.
 - o **Error**. Notifications are sent for Error and Critical events only.
 - o **Warning**. Notifications are sent for Warning, Error, and Critical events only.
 - o **Informational/Resolved**. Notifications are sent for Informational, Warning, Error, Critical, and Resolved events.
2. 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: " < >
3. 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: " ' < >
4. If SNMP notification is enabled, in the Trap Host Address fields enter the IP addresses of hosts that are configured to receive SNMP traps.

To change email notification settings

1. If the mail server is not on the local network, make sure that the gateway IP address was set in [“Using the Configuration Wizard: Network configuration” \(page 38\)](#).
2. Select the **Email** tab.
3. In the SMTP Server address field, enter the IP address of the SMTP mail server to use for the email messages.
4. 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.
5. 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.

6. Perform one of the following:
 - 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.
7. If email notification is enabled, select one of the following Notification Level options:
 - o **Critical**. Notifications are sent for Critical events only.
 - o **Critical, Error**. Notifications are sent for Critical and Error events only.
 - o **Critical, Error, Warning**. Notifications are sent for Critical, Error, and Warning events only.
 - o **Critical, Error, Warning, Resolved**. Notifications are sent for Critical, Error, Warning, and Resolved events.
 - o **Critical, Error, Warning, Resolved, Informational**. Notifications are sent for Critical, Error, Warning, Resolved, and Informational events.
8. 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`.

To change managed logs settings

1. Select the **Managed Logs** tab.
2. Perform one of the following:
 - o To enable managed logs, select the **Enable Managed Logs** check box.
 - o To disable managed logs, clear the **Enable Managed Logs** check box.
3. 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`.
4. Perform one of the following:
 - 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.
5. Click **Next** to proceed to the next step.

Using the Configuration Wizard: Port configuration

To enable the system to communicate with hosts having FC or iSCSI interfaces, you must configure the system's host-interface options. If the current settings are correct, port configuration is optional.

NOTE: The Port configuration panel does not appear for systems with SAS controller modules since it does not have SAS host-interface configuration options.

For systems with a 4-port FC/iSCSI controller module, host ports can be configured as a combination of FC and iSCSI ports. 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 or 10-Gbit/s SFPs. iSCSI port speeds are auto-negotiated.

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

To configure FC ports

1. 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.
2. The FC Connection Mode can be point-to-point or auto:
 - o **point-to-point:** Fibre Channel point-to-point.
 - o **auto:** Automatically sets the mode based on the detected connection type.
3. Click **Next** to proceed to the next step.

To configure iSCSI ports

1. 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 2: 10.10.10.100
 - Controller A port 3: 10.11.10.120
 - Controller B port 2: 10.10.10.110
 - Controller B port 3: 10.11.10.130
 - o Netmask. For IPv4, subnet mask for assigned port IP address.
 - o Gateway. For IPv4, gateway IP address for assigned port IP address.
 - o Default Router. For IPv6, default router for assigned port IP address.
2. 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.

NOTE: CHAP records for iSCSI login authentication must be defined if CHAP is enabled. To create CHAP records, see [“Configuring CHAP” \(page 70\)](#).

- o Link Speed.
 - auto—Auto-negotiates the proper speed.
 - 1 Gbit/s—Forces the speed to 1 Gbit/sec, overriding a downshift that can occur during auto-negotiation with 1-Gbit/sec HBAs. This setting does not apply to 10-Gbit/sec HBAs.
- 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.

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.
- o Enable iSNS. Enables or disables registration with a specified Internet Storage Name Service server, which provides name-to-IP-address mapping.

- iSNS Address. Specifies the IP address of an iSNS server.
- Alternate iSNS Address. Specifies the IP address of an alternate iSNS server, which can be on a different subnet.

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

3. Click **Next** to proceed to the next step.

Using the Configuration Wizard: Confirm the configuration changes

NOTE: For systems with SAS controller modules, this panel appears after the Configure event notification panel since the Port configuration panel is skipped.

Confirm that the changes listed in the wizard panel are correct.

- If they are not correct, click **Previous** to return to previous steps and make necessary changes.
- If they are correct, click **Finish** to apply the settings and finish the wizard. If the changes might disrupt access, confirm the changes.

When processing is complete, you are prompted to add storage. Click **Yes** to open the Add Disk Group panel. Otherwise, click **No**.

Changing system information settings

To change system information settings

1. Perform one of the following:
 - In the Home topic, select **Action > Set System Information**.
 - In the banner, click the system panel and select **Set System Information**.
The Set System Information panel opens.
2. Set the System Name, System Contact person, System Location, and System Information (description) values. The name is shown in the browser title bar or tab. The name, contact, and location 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: " < > \
3. Click **OK**.

Managing 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 standard management interfaces: SMC, CLI, or FTP. 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 [“Changing notification settings” \(page 46\)](#). For information about the MIB, see [“SNMP reference” \(page 121\)](#).

As a user with the `manage` role, you can modify or delete any user other than your current user. Users with the `monitor` role can change all settings for their own user except for user type and role but can only view the settings for other users.

Table 11 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		monitor, manage	WBI, CLI, FTP						
ftp	!ftp		monitor, manage	FTP						

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

User options

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. It cannot already exist in the system or include the following: a space or " , < \
- **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 following options apply only to a standard user:

- **Roles.** Select one or more of the following roles:
 - **Monitor.** Enables the user to view but not change system status and settings. This is enabled by default and cannot be disabled.
 - **Manage.** Enables the user to change system settings.
- **Interfaces.** Select one or more of the following interfaces:
 - **WBI.** Enables access to the SMC.
 - **CLI.** Enables access to the command-line interface.
 - **FTP.** Enables access to the FTP interface, which can be used instead of the SMC 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.
- **Precision Preference.** Select the number of decimal places (1–10) for display of storage-space sizes.
- **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.
 - **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.

- **Timeout.** Select the amount of time that the user's session can be idle before the user is automatically signed out (2–720 minutes).
- **Locale.** Select a display language for the user. Installed language sets include Arabic, Chinese-Simplified, Chinese-Traditional, Dutch, English, French, German, Italian, Japanese, Korean, Portuguese, Russian, and Spanish. The locale determines the character used for the decimal (radix) point, as shown in [“Size representations” \(page 16\)](#).

NOTE: Arabic, Portuguese, and Russian are not supported in this release.

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.
 - **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. 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. 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 IP address of the host system that will receive SNMP traps.

Adding, modifying, and deleting users

To add a new user

1. Perform one of the following:
 - In the Home topic, select **Action > Manage Users**.
 - In the banner, click the user panel and select **Manage Users**.

The User Management panel opens and shows a table of existing users. For information about using tables, see [“Tips for using tables” \(page 14\)](#).
2. Below the table, click **New**.
3. Set the options.
4. Click **Apply**. The user is added and the table is updated.

To create a user from an existing user

1. Perform one of the following:
 - In the Home topic, select **Action > Manage Users**.
 - In the banner, click the user panel and select **Manage users**.

The User Management panel opens and shows a table of existing users. For information about using tables, see [“Tips for using tables” \(page 14\)](#).
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. Click **Apply**. The user is added and the table is updated.

To modify a user

1. Perform one of the following:
 - o In the Home topic, select **Action > Manage Users**.
 - o In the banner, click the user panel and select **Manage users**.
The User Management panel opens and shows a table of existing users. For information about using tables, see [“Tips for using tables” \(page 14\)](#).
2. Select the user to modify.
3. Change the settings. You cannot change the user name. Users with the `monitor` role can change their own settings except for their role and interface settings.
4. Click **Apply**. A confirmation panel appears.
5. Click **Yes** to continue. Otherwise, click **No**. If you clicked Yes, the user is modified.

To delete a user (other than your current user)

1. Perform one of the following:
 - o In the Home topic, select **Action > Manage Users**.
 - o In the banner, click the user panel and select **Manage users**.
The User Management panel opens and shows a table of existing users. For information about using tables, see [“Tips for using tables” \(page 14\)](#).
2. Select the user to delete.
3. Click **Delete**. A confirmation panel appears.
4. Click **Remove** to continue. Otherwise, click **Cancel**. If you clicked Remove, the user is removed and the table is updated.

Changing notification settings

You can enable the system to:

- Send notifications to SNMP trap hosts.
- Send notifications to email addresses when events occur in the system.
- Use the managed logs feature, which transfers log data to a log-collection system. For more information about the managed logs feature, see [“About managed logs” \(page 31\)](#).
- Set remote syslog notification options 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.

To change notification settings

1. Open the Notification Settings panel through either method:
 - o In the footer, click the events panel and select **Set Up Notifications**.
 - o In the Home topic, select **Action > Set Up Notifications**.
2. Change SNMP, email, managed logs settings, and syslog notification settings as described in the first four procedures below.
3. Test the notification settings, as described in [“To test notification settings” \(page 48\)](#).

To change SNMP notification settings

1. Select the **SNMP** tab.
2. If a message near the top of the panel informs you that the SNMP service is disabled, enable it, as described in [“Changing system services settings” \(page 53\)](#).
3. Select one of the following Notification Level options:
 - o **none (disabled)**. All events are excluded from trap notification and traps are disabled.
However, Critical events and managed-logs events are sent regardless of the notification setting.

- **Critical.** Notifications are sent for Critical events only.
 - **Error.** Notifications are sent for Error and Critical events only.
 - **Warning.** Notifications are sent for Warning, Error, and Critical events only.
 - **Informational/Resolved.** Notifications are sent for Informational, Warning, Error, Critical, and Resolved events.
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: " < >
 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: " < >
 6. If SNMP notification is enabled, in the Trap Host Address fields, enter the IP addresses of hosts that are configured to receive SNMP traps.
 7. Click **Apply**.

To change email notification settings

1. If the mail server is not on the local network, make sure that the gateway IP address is set in the System IP Network Configuration panel, as described in [“Changing network interface settings” \(page 54\)](#).
2. Select the **Email** tab.
3. In the SMTP Server address field, enter the IP address of the SMTP mail server to use for the email messages.
4. 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.
5. 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
6. Set the email notification option:
 - To enable email notifications, select the **Enable Email Notifications** check box. This enables the notification level and email address fields.
 - To disable email notifications, clear the **Enable Email Notifications** check box. This disables the notification level and email address fields.
7. If email notification is enabled, select one of the following Notification Level options:
 - **Critical.** Notifications are sent for Critical events only.
 - **Critical, Error.** Notifications are sent for Critical and Error events only.
 - **Critical, Error, Warning.** Notifications are sent for Critical, Error, and Warning events only.
 - **Critical, Error, Warning, Resolved.** Notifications are sent for Critical, Error, Warning, and Resolved events.
 - **Critical, Error, Warning, Resolved, Informational.** Notifications are sent for Critical, Error, Warning, Resolved, and Informational events.
8. 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.
9. Click **Apply**.

To change managed logs settings

1. Select the **Email** tab and ensure that the SMTP Server Address and Sender Domain options are set, as described above.
2. Select the **Managed Logs** tab.

3. 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.
4. 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.
5. Select one of the following options:
 - o To use the 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 the pull mode, clear the **Include logs as an email attachment** check box.
6. Click **Apply**.

To configure syslog notification of events

1. Select the **Syslog** tab.
2. 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); 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.
3. Click **Next** to proceed to the next step.

To test notification settings

1. Click **Send Test Event**. A test notification is sent to each configured trap host and email address.
2. Verify that the test notification reached each configured trap host and email address.
3. If the managed logs option is enabled, click **Send Log Test**. A test notification is sent to the log-collection system.
4. Verify that the test notification reached the log-collection system.

Managing scheduled tasks

You can modify or delete scheduled tasks to:

- Create snapshots
- Reset snapshots
- Run replications

To modify a schedule

1. In the Home topic, select **Action > Manage Schedules**. The Manage Schedules panel opens.
2. Select the schedule to modify. The schedule's settings appear at the bottom of the panel.
3. Modify the settings.
4. Click **Apply**. A confirmation panel appears.
5. Click **Yes** to continue. Otherwise, click **No**. If you clicked Yes, the schedule is modified.
6. Click **OK**.

To delete a schedule

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 **Apply** to continue. Otherwise, click **No**. If you clicked Apply, the schedule was deleted.
5. Click **OK**.

Installing a license

A license is required to use the Performance tier, expand the maximum number of snapshots, and use the replication feature and VSS.

NOTE: VSS is not supported in this release.

Viewing the status of licensed features

1. In the Home topic, select **Action > Install License**. The License Settings panel opens and shows the following information about each licensed feature:
 - Feature. The feature name.
 - Base. One of the following:
 - The number of standard snapshots that users can create without a license.
 - N/A. Not applicable.
 - 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.
 - In Use. One of the following:
 - The number of standard snapshots that exist.
 - N/A. Not applicable.
 - Max Licensable. One of the following:
 - The number of standard snapshots that the maximum license supports.
 - N/A. Not applicable.
 - 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 permanent license

1. Verify the following:
 - The license file is saved to a network location that you can access from the SMC.
 - You are signed into the controller enclosure for which the file is generated.
2. In the Home topic, select **Action > Install License**. The License Settings panel opens.
3. On the Permanent License tab, click **Choose File** to locate and select the license file.
4. Click **OK**. The license settings table is updated and, for each feature included in the license, the *Expiration* value changes to *Never*.

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. If installed disks are part of a virtual disk group or are global spares, unique color codes identify them as such. For information on the specific colors used, see [“Color codes” \(page 15\)](#).

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 any of the components, 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, 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, FDE lock key, 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.
- *FDE State* refers to the FDE state of the disk. For more information about FDE states, see the CLI Reference Guide.
- *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:

NOTE: FC, iSCSI, and SAS host port information displays only when the host port is installed on the system.

Table 12 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

Table 12 Additional information for rear view of enclosure

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.

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” \(page 14\)](#).

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, 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.
 - SPARE. The disk is configured as a 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 because the system is secured, or the disk is locked to data access, or the disk is from an unsupported vendor.
- 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.

Managing global spares

If a disk in any fault-tolerant virtual 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 for virtual disk groups). You can designate a maximum of 16 global spares, but at least one disk group must exist before you can add a global spare. (To designate up to 64 global spares, use the CLI.) 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 Change Global Spares panel contains a single disk set, which consists of the disks selected as global spares. The Disk Sets summary includes a row that shows the total space of the disk set and amount of space allocated for spares. Also, a table is located below the row that contains several fields, including the Disks and Size fields. The total space and Size field values reflect the cumulative amount of storage for the selected disks. The Disks field shows the number of spares selected.

Underneath the Disk Sets summary are one or more disk tables depending on how many enclosures your system has. Each table represents an enclosure and all of its disks. To see more information about an enclosure or disks, such as the disk type, capacity, and sector format, hover the cursor over an enclosure or disk. The Enclosure Information or Disk Information panel appears. [“Viewing pools” \(page 72\)](#) contains more details about the Disk Information panel.

If installed disks are part of a virtual disk group or are global spares, unique color codes identify them as such. For information on the specific colors used, see [“Color codes” \(page 15\)](#).

NOTE: Disk groups support a mix of 512n and 512e disks. 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” \(page 18\)](#).

To change the system's global spares

1. In the System topic, select **Action > Change Global Spares**. The Change Global Spares panel shows information about available disks in the system. Existing spares are labeled GLOBAL SP.
 - o In the Disk Sets summary, the number of white slots in the Disks field shows how many spares you can add. The colored slots show how many disks you have selected to become spares or have already added as spares.
 - o In each disk table, which visually represents the disks for an enclosure, only existing global spares and suitable available disks are selectable.
2. Select spares to remove, disks to add as spares, or both.
3. Click **Change**. If the task succeeds, the panel is updated to show which disks are now global spares.

Changing system services settings

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. To allow specific users to access the SMC (the web browser interface), CLI, or FTP interfaces, see [“Managing users” \(page 43\)](#).

To change system services settings

1. Perform one of the following:
 - o In the banner, click the system panel and select **Set Up System Services**.
 - o In the System topic, select **Action > Set Up System Services**.The System Services panel opens.
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. You can enable use of **HTTP** and/or **HTTPS** for increased security. If you disable both, you will lose access to this interface.
 - o Command Line Interface (CLI). An advanced-user interface that is used to write scripts to manage the system. You can enable use of **SSH** (secure shell) for increased security, **Telnet**, or both.

- 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.
- File Transfer Protocol (FTP). A secondary interface for installing firmware updates, downloading logs, and installing a license.
- Simple Network Management Protocol (SNMP). Used for remote monitoring of the system through your network.
- Service Debug. Used for technical support only. Enables or disables debug capabilities, including Telnet debug ports and privileged diagnostic user IDs.

NOTE: Properly shut down the debug console by using the CLI command `set protocols debug disable`. Do not just close the console directly or by using the CLI command `exit`.

- 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.
 - 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.
3. Click **OK**. If any unsecure interfaces are enabled, a confirmation panel will appear.
 4. Click **Yes** to confirm use of unsecure interfaces. Otherwise, click **No**.

Changing network interface settings

You can change addressing parameters for the network port in each controller module. You can set static IP values or use DHCP. When setting static IP values, you can use IPv4 format.


In 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, current addressing is unchanged. You must have some means of determining what addresses have been assigned, such as the list of bindings on the DHCP server.

Each controller has the following factory-default IP settings:

- IP address source: DHCP
- 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

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: 0.0.0.0

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

To use DHCP to obtain IP values for network ports

1. In the System topic, select **Action > Set Up Network**. The System IP Network Configuration panel opens.

2. Set IP address source to **DHCP** and click **OK**. If the controllers successfully obtain IP values from the DHCP server, the new IP values appear.
3. Record the new addresses.
4. Sign out and try to access the SMC using the new IP addresses.

To set static IP values for network ports

1. Determine the IP address, subnet mask, and gateway values to use for each network port.
2. In the System topic, select **Action > Set Up Network**. The System IP Network Configuration panel opens.
3. Set IP address source to **manual**.
4. Enter IP address, subnet mask, and gateway values for each controller. You must set a unique IP address for each controller.

NOTE: The following IP addresses are reserved for internal use by the storage system: 192.168.200.253, 192.168.200.254, 172.22.255.253, 172.22.255.254, and 127.0.0.1.

5. Record the IP values you assign.
6. Click **OK**.
7. Sign out and try to access the SMC using the new IP addresses.

Changing host port settings

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

For a system with a 4-port SAS controller module, there are no host-interface options.

For a system with a 2-port SAS controller module, host ports can be configured to use standard cables. A standard cable can connect one port on a SAS host to one controller port, using four PHY lanes per port. 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.

CNC host ports can be configured as a combination of FC or iSCSI ports. 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 or 10-Gbit/s SFPs. iSCSI port speeds are auto-negotiated.

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

To configure FC ports

1. In the System topic, select **Action > Set Up Host Ports**. The Host Ports Settings panel opens.
2. 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.
3. The FC Connection Mode can be point-to-point or auto:
 - o **point-to-point:** Fibre Channel point-to-point.
 - o **auto:** Automatically sets the mode based on the detected connection type.
4. Click **OK**.
5. Click **Apply**. Otherwise, click **Cancel**. If you clicked Apply, the ports are configured.
6. Click **OK**.

To configure iSCSI ports


1. In the System topic, select **Action > Set Up Host Ports**. The Host Ports Settings panel opens.
2. 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 2: 10.10.10.100
 - Controller A port 3: 10.11.10.120
 - Controller B port 2: 10.10.10.110
 - Controller B port 3: 10.11.10.130
 - o Netmask. For IPv4, subnet mask for assigned port IP address.
 - o Gateway. For IPv4, gateway IP address for assigned port IP address.
 - o Default Router. For IPv6, default router for assigned port IP address.
3. 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.

NOTE: CHAP records for iSCSI login authentication must be defined if CHAP is enabled. To create CHAP records, see [“Configuring CHAP” \(page 70\)](#).

- o Link Speed.
 - auto—Auto-negotiates the proper speed.
 - 1 Gbit/s—Forces the speed to 1 Gbit/sec, overriding a downshift that can occur during auto-negotiation with 1-Gbit/sec HBAs. This setting does not apply to 10-Gbit/sec HBAs.
- 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.

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.
- o Enable iSNS. Enables or disables registration with a specified Internet Storage Name Service server, which provides name-to-IP-address mapping.
- o iSNS Address. Specifies the IP address of an iSNS server.
- o Alternate iSNS Address. Specifies the IP address of an alternate iSNS server, which can be on a different subnet.

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

4. Click **OK**.
5. Click **Apply**. Otherwise, click **Cancel**. If you clicked Apply, the ports are configured.
6. Click **OK**.

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 (either the first and second ports or the third and fourth ports). 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 Hardware Installation and Maintenance 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. 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” \(page 30\)](#). For information about how to enable PFU using the `set advanced-settings` CLI command, see the CLI Reference Guide.

To monitor the progress of a firmware-update operation by using the activity progress interface, see [“Using the activity progress interface” \(page 61\)](#).

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” \(page 116\)](#).
- 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 the firmware in each controller module by loading a firmware file obtained from the Lenovo web download site at support.lenovo.com.

To prepare to update controller module firmware

1. Follow the best practices in “[Best practices for firmware update](#)” (page 58).
2. Obtain the appropriate firmware file and download it to your computer or network.
3. If the storage system has a single controller, stop I/O to the storage system before you start the firmware update.

To update controller module firmware

1. Perform one of the following:
 - In the banner, click the system panel and select **Update Firmware**.
 - In the System topic, select **Action > Update 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. 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 a 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 10–20 minutes for the partner controller to be updated.

4. Clear your web browser cache, then sign in to the SMC. 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 through the `partner-firmware-upgrade` parameter of the `set advanced-settings` CLI command, 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 can contain one or 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 the firmware in each expansion module by loading a firmware file available only from Lenovo at support.lenovo.com and follow the steps below.

To prepare to update expansion module firmware

1. Follow the best practices in [“Best practices for firmware update” \(page 58\)](#).
2. Obtain the appropriate firmware file and download it to your computer or network.
3. If the storage system has a single controller, stop I/O to the storage system before starting the firmware update.

To update expansion module firmware

1. Perform one of the following:
 - o In the banner, click the system panel and select **Update Firmware**.
 - o In the System topic, select **Action > Update 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 3 minutes to update each EMP in an expansion enclosure. Wait for a message that the code load has completed.

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

Updating disk-drive firmware

You can update disk-drive firmware by loading a firmware file obtained from your reseller.

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” \(page 58\)](#).
2. Obtain the appropriate firmware file and download it to your computer or network.
3. Read documentation from the disk-drive manufacturer to determine whether the disk drives must be power cycled after firmware update.
4. 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. Perform one of the following:
 - o In the banner, click the system panel and select **Update Firmware**.
 - o In the System topic, select **Action > Update 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 [“Changing system services settings” \(page 53\)](#).
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:
 - o *controller-address* is required and specifies the IP address of a controller network port.
 - o *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.
 - o *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 13 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.

Table 13 Activity progress properties and values

Property	Value
Status	<p>The status of the component representing its progress/completion state.</p> <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).
Code	<p>A numeric code indicating the status.</p> <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

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:
 - o To secure the system now, click the **Set** checkbox, then select **Set and Secure**. A dialog box will confirm the passphrase was changed successfully.
 - o 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” \(page 63\)](#).

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. Disks will remain in the secured, unlocked 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 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” \(page 62\)](#).

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. In the Secure System section, click the **Secure** button.
4. Click **Clear**. A dialog box displays.
5. Perform one of the following:
 - o To secure the system, click **Set and Secure**.
 - o To cancel the request, click **Cancel**.

Securing the system

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

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” \(page 62\)](#).

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 dialog box displays.
4. Perform one of the following:
 - o To secure the system, click **Yes**.
 - o To cancel the request, click **No**.

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” \(page 77\)](#). 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 **Yes**.
 - o To cancel the request, click **No**.

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.

 **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.

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 SMC 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” \(page 110\)](#).


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 **Yes** to continue. Otherwise, click **No**. If you clicked **Yes**, 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.

 **CAUTION:** You can continue to use the CLI when either or both Storage Controllers are shut down, but information shown might be invalid.

To perform a shut down

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 **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 **Yes** to continue. Otherwise, click **No**. If you clicked **Yes**, 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” \(page 14\)](#). For more information about hosts, see [“About initiators, hosts, and host groups” \(page 26\)](#). The Hosts topic also enables users to map initiators (see [page 92](#)) and view map details (see [page 95](#)).

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 a manually created initiator.
- 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” \(page 95\)](#).

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 the nickname of any initiator.

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**. The initiators are deleted 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 have the same mappings as all other initiators in the host. This means that the initiator must be mapped with the same access, port, and LUN settings to the same volumes or volume groups.

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 have the same mappings as all other members of the 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” \(page 70\)](#).

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” \(page 55\)](#).

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” \(page 102\)](#).

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” \(page 14\)](#). For more information about pools, see [“About pools” \(page 22\)](#). For more information about disk groups, see [“About disk groups” \(page 18\)](#).

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	Virtual: 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” \(page 78\)](#).

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.
 - 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, 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, FDE lock key, 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.
- *FDE State* refers to the FDE state of the disk. For more information about FDE states, see the CLI Reference Guide.
- *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 type 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 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” \(page 18\)](#).

Add Disk Group panel overview

There are three sections that comprise the Add Disk Group panel. The top section provides options for the disk group, such as the type, name, and RAID level of the disk group. The options that appear vary depending on the type of disk group selected.

The middle section contains the disk selection sets summary, which presents cumulative data for the disks selected for the disk group. The amount of disk space (total, available, and overhead) appears, as do the RAID and disk types that have been selected for the disk group.

The summary also contains the Disks bar, which shows the number of disks selected, and the **Complete** check box. The Disks bar appears for disks intended for use in a RAID configuration or in a read-cache disk group. The **Complete** check box indicates if the minimum number of disks needed for the configuration have been selected. It automatically changes from ☐ to ☒ when the minimum has been selected. The options that appear in the middle section vary depending on the type of disk group selected.

In the bottom section are one or more disk tables depending on the number of enclosures that your system has. Each table represents an enclosure and all of its disks. Open check boxes appear on available disks. To see more information about an enclosure or disks, such as the disk type, capacity, and sector format, hover the cursor over an enclosure ear or disk. The Enclosure Information or Disk Information panel appears. [“Viewing pools” \(page 72\)](#) contains more details about the Disk Information panel.

If installed disks are part of a virtual disk group or are global spares, unique color codes identify them as such. For information on the specific colors used, see [“Color codes” \(page 15\)](#).

Using SSDs in a disk group

For the best performance for virtual disk groups with SSDs, use at least a minimum of four SSDs with one mirrored pair of drives (RAID 1) per controller. RAID 5 and RAID 6 are also preferable for SSDs. However, they require more drives if you are following the best practice of having one disk group owned by each controller. In this configuration, six SSDs for RAID 5 and eight SSDs for RAID 6 are required. For data integrity, use a fault-tolerant RAID level for a disk group of SSDs.

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

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 modify it. If your organization's needs change, you can modify your storage amount by adding new virtual disk groups or deleting 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 in virtual disk groups, which automatically use the Performance tier for SSDs. 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 groups in the same tier within a virtual pool should have the same 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 have dequarantined the existing disk group. For information on quarantining and dequarantining disk groups, see the CLI documentation.

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: " , < \

- RAID Level. Select one of the following RAID levels when creating a virtual disk group:
 - **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 automatically 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” \(page 72\)](#) 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. Click **Yes** to continue. Otherwise, click **No**. If you clicked Yes, the disk group modification has started.
6. To close the confirmation panel, click **OK**.

Removing disk groups

You can remove disk groups. It is possible to 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.

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.

When removing a disk group in an all-flash array (a storage system using only SSDs), the system will not drain data from one virtual disk group to another one because similar wear across the SSDs is likely, so more failures may be imminent.

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” \(page 72\)](#) 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**.
4. Click **Yes** to continue. Otherwise, click **No**. 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.

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. Then, select a 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” \(page 72\)](#) 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” \(page 84\)](#).

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 25%.
- **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 50%. 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” \(page 25\)](#).

NOTE: 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. 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.

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” \(page 72\)](#) 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 the storage system was operating outside the normal temperature range, 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 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.

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, see the `set advanced-settings` command in the CLI Reference Guide.

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” \(page 72\)](#) 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 background scrub option automatically checks all disk groups for disk defects. For details, see documentation for the CLI `set advanced-settings` command. 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, 6; mirror mismatches for RAID 1 and 10; and media errors for all RAID levels.

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 "foreground" scrub performed by Scrub Disk Group is typically faster than a background 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 using the `abort scrub` CLI command.

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” \(page 96\)](#). For more information about snapshots, see [“About snapshots” \(page 28\)](#). For information about using tables, see [“Tips for using tables” \(page 14\)](#).

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, write policy, optimization, read-ahead size, tier affinity, health
--------------------	--

For more information about write policy and read-ahead size, see [“Modifying a volume” \(page 85\)](#).

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, 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).
- Type. Shows the following snapshot type:
 - Standard snapshot. Snapshot of a standard volume.

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” \(page 95\)](#).

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, primary volume, primary volume serial, secondary volume, secondary volume serial, peer connection, associated schedule name, current run progress, current run start time, current run estimated time to completion, current run transferred date, last successful run, last run start time, last run end time, last run transferred date, 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"> ◦ Reset snapshot - snapshot name, snapshot serial ◦ Take snapshot - source volume, source volume serial, prefix, count, last created ◦ Volume copy - source volume, source volume serial, destination disk group, destination disk group serial, prefix, 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 copies of a virtual volume with the same settings. In the latter case, the copies will have the same base name with a numeric suffix (starting at 0000) to make each name unique. 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” \(page 81\)](#) 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 SMC 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. See the system configuration limits topic in SMC 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” \(page 26\)](#).

To see more information about a volume, hover the cursor over the volume in the table. [“Viewing volumes” \(page 81\)](#) 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**.
6. Optional: In the Write Optimization list, select **Standard** or **No-mirror**.
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).
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**.

Adding volumes to a volume group

You can add virtual volumes to a new or existing virtual volume group.

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, port, and LUN 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.

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” \(page 86\)](#).

To see more information about a volume, hover the cursor over the volume in the table. [“Viewing volumes” \(page 81\)](#) 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” \(page 81\)](#) contains more details about the Volume Information panel that appears, including how to view volumes and volume groups that are part of a replications 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” \(page 81\)](#) 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 1–32 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 1–32 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

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

For virtual snapshots, 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 virtual 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 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.

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” \(page 81\)](#) contains more details about the Volume Information panel that appears.

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**. 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 81) 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. You can create more snapshots with an upgrade license. You can create snapshots immediately or schedule snapshot creation.

A base of 128 snapshots is included with all systems without an additional license.

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 81) contains more details about the Volume Information and Snapshot Information panels that appear.

To create virtual snapshots

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

NOTE: You can also select a combination of virtual 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 1-32 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” \(page 48\)](#).

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 action is not allowed for a replication snapshot.

For virtual snapshots, 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” \(page 81\)](#) 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:
 - o Select the **Schedule** check box.

- 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:
 - 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**:
 - o If the **Schedule** check box was not selected, the snapshot is created. You can remount the snapshot.
 - o If **Schedule** is selected, the schedule is created and can be viewed in the Manage Schedules panel, as described in [“Managing scheduled tasks” \(page 48\)](#). 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 SMC—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.

To create a replication set

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.
 - o To replicate the selected volume only, select **Single Volume**. This option is the default.
 - o 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: Select the **Scheduled** check box to schedule recurring replications.
9. Click **OK**.
10. 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 [“Scheduling replications” \(page 107\)](#).
 - 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 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. You can initiate replications from a replication set's primary system only. For information on scheduling replications, see [“Scheduling replications” \(page 107\)](#).

You can initiate a replication from both the Replications and Volumes topics. For information on how to initiate a replication, see [“Initiating replication” \(page 106\)](#).

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.

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” \(page 14\)](#).

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” \(page 95\)](#).

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. This behavior is expected by some operating systems, such as Microsoft Windows, which can immediately discover the volume. 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 initiators, you can create an *explicit mapping*. An explicit mapping can use different access mode, LUN, and port settings to allow or prevent access by an initiator 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 SMC. It will be replaced in the SMC 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: 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:

- In the Hosts topic, select the initiators to map and select **Action > Map Initiators**.
- In the Volumes topic, select the volumes to map and select **Action > Map Volumes**.
- In the Mapping topic, select **Map** to create a new mapping.
- 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 14 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 15 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 can only be 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 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.
- To save a new mapping or edits to an existing mapping, in the Action column, select **Save**.
- 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**.

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:
 - In the Hosts or Volumes topic, in the Related Maps table, select at least one mapping.
 - 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” \(page 14\)](#).
 - 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.) This feature 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 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” \(page 102\)](#), [“Deleting a peer connection” \(page 104\)](#), and [“Modifying a peer connection” \(page 104\)](#).

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 two internal snapshots for each volume, which are created on both the primary and secondary systems. These snapshots are used internally for replication operations.

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. These snapshots also can be used in disaster recovery. For more information on replication sets, see [“Creating a replication set from the Replications topic” \(page 105\)](#), [“Creating a replication set from the Volumes topic” \(page 90\)](#), [“Modifying a replication set” \(page 106\)](#), and [“Deleting a replication set” \(page 106\)](#).

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” \(page 108\)](#), [“Initiating replication” \(page 106\)](#), [“Initiating replication from the Volumes topic” \(page 91\)](#), [“Resuming a replication” \(page 109\)](#), and [“Suspending a replication” \(page 108\)](#).

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, see [“Scheduling replications” \(page 107\)](#).

Initial replication

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

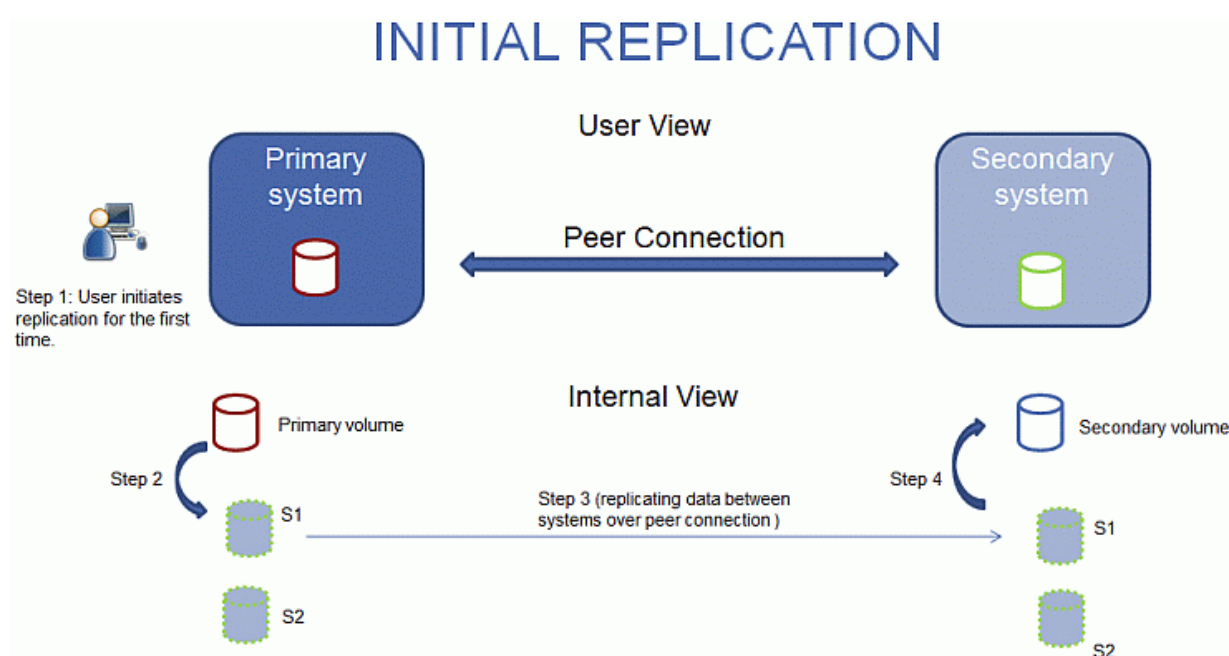


Figure 1 Replication process for initial replication

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. 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.

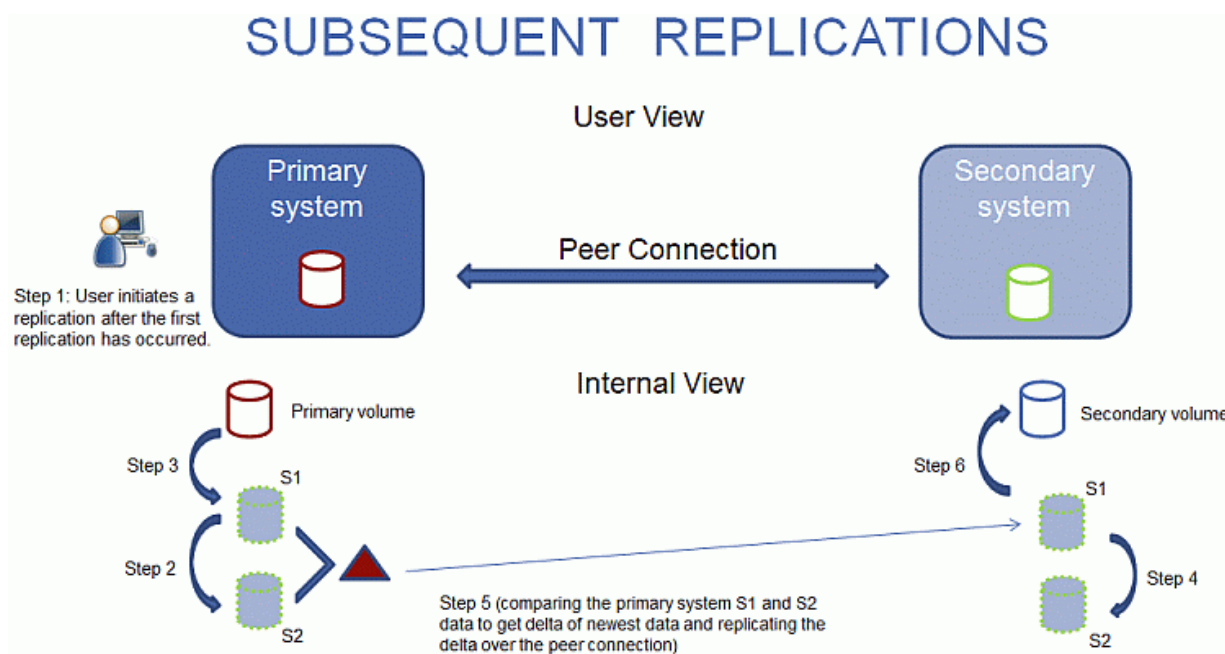


Figure 2 Replication process for replications subsequent to the initial replication.

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 do not count against the maximum number of licensed snapshots for the system.

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.

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 the replication feature 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.

Disaster recovery

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

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

Accessing the data while keeping the replication set intact

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

1. Create a snapshot of the secondary volume, 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 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 and replication sets that are defined in the system. For information about using tables, see [“Tips for using tables” \(page 14\)](#). For more information about replication, see [“About replicating virtual volumes” \(page 96\)](#).

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: 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.

- 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; primary volume or volume group name and serial number; secondary volume or volume group name and serial number; peer connection name; replication schedule name, current run progress, current run start time, current run estimated time to completion, current run transferred data, last success time, last run start time, last run end time, last run transferred data, last run status, and last run error status
------------------	---

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 an IP address of a single available iSCSI host port on the remote system. Only iSCSI host ports are used for the peer connection. FC or SAS 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 virtual replication.
- Both systems must have iSCSI ports.
- Each system must have a virtual pool.
- Neither system can have a linear replication set.
- 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.

The limit is one peer connection per storage system.

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 the replication feature 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 address (IP) for the remote system.
4. Click **OK**.
5. 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**.

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.

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.

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. Enter a new name for the peer connection or a new address (IP) for the remote system. (You cannot change both.) The 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**. 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” \(page 106\)](#).

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.

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 SMC—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.

To create a replication set

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: Select the **Scheduled** check box to schedule recurring replications.
9. Click **OK**.

10. 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 [“Scheduling replications” \(page 107\)](#).
- 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 the name of a replication set. Volume membership of a replication cannot change for the life of the replication set.

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**.
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. Click **OK**. The name of the replication set is deleted and the Replications Sets table is updated.

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” \(page 108\)](#) or [“Suspending a replication” \(page 108\)](#).

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 replication

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. You can initiate replications from a replication set's primary system only. For information on scheduling replications, see [“Scheduling replications” \(page 107\)](#). You can initiate a replication from both the Replications and Volumes topics.

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.
-

To manually initiate replication

1. In either the Replications or Volumes topic, select a replication set in the Replication Sets table.
2. Select **Action > Initiate Replication**.
3. Click **OK**. 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**.

Scheduling replications

After you have created a replication set, you can schedule replication. You can schedule replications from a replication set's primary system only. For more information on replications, see [“Initiating replication” \(page 106\)](#).

NOTE: Virtual storage replication tasks are not queued. If a replication task is running and the time comes for that replication task to start again, that task will be skipped, though it will be counted against the schedule's count constraint (if set). Instead of an open-ended schedule, you can specify the number of replications that can occur for the replication task as part of the procedure for scheduling a replication.

To schedule a replication

1. In the Replications topic, select a replication set in the Replication Sets table.
2. Select **Action > Schedule Replications**. The Schedule Replications panel opens.
3. Select the **Schedule** check box.
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). The minimum interval is one hour.
5. Optional: If you want the task to run more than once, select the **Schedule** check box.
 - Specify how often the task should repeat. Enter a number and select the appropriate time unit. Replications can recur no less than 60 minutes apart.
 - 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.
 - 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.
 - 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 **OK**. The schedule is created.

To manage a replication schedule

1. In the Replications topic, select a replication set in the Replication Sets table.
2. Select **Action > Manage Schedules**.
3. Set the options:
 - o 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).
4. 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 60 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.
5. Click **OK**. The schedule is modified.

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

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” \(page 108\)](#) or [“Resuming a replication” \(page 109\)](#). If replications are attempted during the suspended period (including scheduled replications), the replications will fail.

To suspend a replication

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

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*.


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” \(page 30\)](#).

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” \(page 14\)](#).
 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” \(page 110\)](#).
 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 16 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 16 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 16 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” \(page 30\)](#).

You can update historical statistics.

To update displayed historical statistics

1. Display a historical statistics graph as described in [“Viewing performance statistics” \(page 110\)](#).
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 SMC 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 SMC 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 SMC 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 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) and the Lenovo serial number.

The  icon indicates that the panel has a menu. Click anywhere in the panel to display a menu to change system information settings ([page 43](#)) and system services settings ([page 53](#)), update firmware ([page 58](#)), restart or shut down controllers ([page 65](#)) and view SSL certificate information ([page 114](#)).

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, or you have selected the tab for a partner controller in a single-controller system.

You can use your own certificates by uploading them through FTP 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 65).

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 SMC and the storage system. The connection information table shows the icon that displays for each state.

Table 17 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*.

Hover the cursor over this panel to display the System Date/Time panel, which shows NTP settings.

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, such as for Daylight Saving Time. You must make such adjustments 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 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” \(page 43\)](#).

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 46](#)), save log data ([page 116](#)), and view system information ([page 50](#)).

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 SMC, 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” \(page 118\)](#) and change notification settings on [“Changing notification settings” \(page 46\)](#).

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 [“Changing notification settings” \(page 46\)](#). For information about using tables, see [“Tips for using tables” \(page 14\)](#).

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 Hardware Installation and Maintenance Guide
- The topics about verifying component failure in your product's Hardware Installation and Maintenance Guide
- 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” \(page 15\)](#).

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 virtual 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 that the data written to the pool takes. 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 virtual disk groups, both total and by pool
- Reserved. The reserved space for virtual disk groups, both total and by pool
- Allocated. The allocated space for virtual disk groups, both total and by pool
- Unallocated. The unallocated space for virtual 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” \(page 15\)](#).

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” \(page 15\)](#).

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 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.

To view past notifications for this SMC session, click the activity panel in the footer and select Notification History. For more information, see [“Viewing the notification history” \(page 120\)](#).

Viewing the notification history

The Notification History panel shows past activity notifications for this SMC 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.

A Other management interfaces

SNMP reference

This appendix describes the Simple Network Management Protocol (SNMP) capabilities that Lenovo ThinkSystem 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 Lenovo ThinkSystem storage systems is 1.3.6.1.4.1.19046. 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, `1straps.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 SMC.

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 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 SMC. The statistics section is not implemented.

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

Table 18 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 18 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 18 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 SMC.	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 127)
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 127)
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 128)
connUnitSensorName	Display string containing textual identification of the sensor intended primarily for operator use	See “External details for connUnitSensorTable” (page 128)
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 18 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 128)
connUnitSensorCharacteristic	Characteristics being monitored by this sensor	See “External details for connUnitSensorTable” (page 128)
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 18 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 129)
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 18 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

External details for certain FA MIB 2.2 objects

Tables in this section specify values for certain objects described in [Table 18](#).

External details for connUnitRevsTable

Table 19 connUnitRevsTable index and description values

connUnitRevsIndex	connUnitRevsDescription
1	CPU Type for Storage Controller (Controller A)
2	Bundle revision for Controller (Controller A)
3	Build date for Storage Controller (Controller A)
4	Code revision for Storage Controller (Controller A)
5	Code baselevel for Storage Controller (Controller A)
6	FPGA code revision for Memory Controller (Controller A)
7	Loader code revision for Storage Controller (Controller A)
8	CAPL 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)

Table 19 connUnitRevsTable index and description values (continued)

connUnitRevsIndex	connUnitRevsDescription
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 20 connUnitSensorTable index, name, type, and characteristic values**

connUnitSensorIndex	connUnitSensorName	connUnitSensorType	connUnitSensor Characteristic
1	Onboard Temperature 1 (Controller A)	board(8)	temperature(3)
2	Onboard Temperature 1 (Controller B)	board(8)	temperature(3)
3	Onboard Temperature 2 (Controller A)	board(8)	temperature(3)
4	Onboard Temperature 2 (Controller B)	board(8)	temperature(3)
5	Onboard Temperature 3 (Controller A)	board(8)	temperature(3)
6	Onboard Temperature 3 (Controller B)	board(8)	temperature(3)
7	Disk Controller Temperature (Controller A)	board(8)	temperature(3)
8	Disk Controller Temperature (Controller B)	board(8)	temperature(3)
9	Memory Controller Temperature (Controller A)	board(8)	temperature(3)
10	Memory Controller Temperature (Controller B)	board(8)	temperature(3)
11	Capacitor Pack Voltage (Controller A)	board(8)	power(9)
12	Capacitor Pack Voltage (Controller B)	board(8)	power(9)
13	Capacitor Cell 1 Voltage (Controller A)	board(8)	power(9)
14	Capacitor Cell 1 Voltage (Controller B)	board(8)	power(9)
15	Capacitor Cell 2 Voltage (Controller A)	board(8)	power(9)
16	Capacitor Cell 2 Voltage (Controller B)	board(8)	power(9)
17	Capacitor Cell 3 Voltage (Controller A)	board(8)	power(9)
18	Capacitor Cell 3 Voltage (Controller B)	board(8)	power(9)
19	Capacitor Cell 4 Voltage (Controller A)	board(8)	power(9)

Table 20 connUnitSensorTable index, name, type, and characteristic values (continued)

connUnitSensorIndex	connUnitSensorName	connUnitSensorType	connUnitSensorCharacteristic
20	Capacitor Cell 4 Voltage (Controller B)	board(8)	power(9)
21	Capacitor Charge Percent (Controller A)	board(8)	other(2)
22	Capacitor Charge Percent (Controller B)	board(8)	other(2)
23	Overall Status	enclosure(7)	other(2)
24	Upper IOM Temperature (Controller A)	enclosure(7)	temperature(3)
25	Lower IOM Temperature (Controller B)	enclosure(7)	temperature(3)
26	Power Supply 1 (Left) Temperature	power-supply(5)	temperature(3)
27	Power Supply 2 (Right) Temperature	power-supply(5)	temperature(3)
28	Upper IOM Voltage, 12V (Controller A)	enclosure(7)	power(9)
29	Upper IOM Voltage, 5V (Controller A)	enclosure(7)	power(9)
30	Lower IOM Voltage, 12V (Controller B)	enclosure(7)	power(9)
31	Lower IOM Voltage, 5V (Controller B)	enclosure(7)	power(9)
32	Power Supply 1 (Left) Voltage, 12V	power-supply(5)	power(9)
33	Power Supply 1 (Left) Voltage, 5V	power-supply(5)	power(9)
34	Power Supply 1 (Left) Voltage, 3.3V	power-supply(5)	power(9)
35	Power Supply 2 (Right) Voltage, 12V	power-supply(5)	power(9)
36	Power Supply 2 (Right) Voltage, 5V	power-supply(5)	power(9)
37	Power Supply 2 (Right) Voltage, 3.3V	power-supply(5)	power(9)
38	Upper IOM Voltage, 12V (Controller A)	enclosure(7)	currentValue(6)
39	Lower IOM Voltage, 12V (Controller B)	enclosure(7)	currentValue(6)
40	Power Supply 1 (Left) Current, 12V	power-supply(5)	currentValue(6)
41	Power Supply 1 (Left) Current, 5V	power-supply(5)	currentValue(6)
42	Power Supply 2 (Right) Current, 12V	power-supply(5)	currentValue(6)
43	Power Supply 2 (Right) Current, 5V	power-supply(5)	currentValue(6)

External details for connUnitPortTable**Table 21 connUnitPortTable index and name values**

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

Configuring SNMP event notification in the SMC

1. Verify that the storage system's SNMP service is enabled. See [“Changing system services settings” \(page 53\)](#).
2. Configure and enable SNMP traps. See [“Changing notification settings” \(page 46\)](#).
3. Optionally, configure a user account to receive SNMP traps. See [“Managing users” \(page 43\)](#).

SNMP management

You can manage storage devices using SNMP with a network management system. 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 [“Changing system services settings” \(page 53\)](#). To use SNMPv3, it must be configured in both the storage system and the network management system that intends to access the storage system or receive traps from it. In the storage system, SNMPv3 is configured through the creation and use of SNMP user accounts, as described in [“Managing users” \(page 43\)](#). 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 Lenovo enterprise traps MIB, `lstraps.mib`. This MIB defines the content of the SNMP traps that Lenovo ThinkSystem storage systems generate.

```
-- -----
-- Lenovo Storage Array MIB for SNMP Traps
--
--
-- Copyright 2015 Lenovo
-- All rights reserved. Use is subject to license terms.
--
-- -----

LSTRAPS-MIB
-- Last edit date: April 29th, 2015
DEFINITIONS ::= BEGIN
    IMPORTS
        enterprises
            FROM RFC1155-SMI
        TRAP-TYPE
            FROM RFC-1215
        connUnitEventId, connUnitEventType, connUnitEventDescr
            FROM FCMGMT-MIB;

    --Textual conventions for this MIB

    -----
    lenovo    OBJECT IDENTIFIER ::= { enterprises 19046 }

    -- Related traps

    lsEventInfoTrap TRAP-TYPE
        ENTERPRISE lenovo
        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
::= 1

lsEventWarningTrap TRAP-TYPE
    ENTERPRISE lenovo
    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
::= 2

lsEventErrorTrap TRAP-TYPE
    ENTERPRISE lenovo
    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
::= 3

lsEventCriticalTrap TRAP-TYPE
    ENTERPRISE lenovo
    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
::= 4

lsEventResolvedTrap TRAP-TYPE
ENTERPRISE lenovo
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 "Informational storage event"
--#SUMMARY "Informational storage event # %d, type %d, description: %s"
--#ARGUMENTS {0,1,2}
--#SEVERITY INFORMATIONAL
--#TIMEINDEX 6
::= 5

END

```

Using FTP

Although the SMC is the preferred interface for downloading log data and historical disk-performance statistics, updating firmware, installing a license, and installing a security certificate, you can also use FTP to do these tasks.

-
- ❗ **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 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 client. A GUI-based FTP client might not work.

To download system logs

1. In the SMC, prepare to use FTP:
 - a. Determine the network-port IP addresses of the system's controllers. See ["Changing network interface settings" \(page 54\)](#).
 - b. Verify that the system's FTP service is enabled. See ["Changing system services settings" \(page 53\)](#).
 - c. Verify that the user you will log in as has permission to use the FTP interface. See ["To modify a user" \(page 46\)](#).
2. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the destination directory for the log file.
3. Enter:


```
ftp controller-network-address
```

 For example:


```
ftp 10.1.0.9
```
4. Log in as a user that has permission to use the FTP interface.
5. Enter:


```
get logs filename.zip
```

 where *filename* is the file that will contain the logs. It is recommended to choose a filename that identifies the system, controller, and date.
 For example:


```
get logs Storage2_A_20120126.zip
```

 Wait for the message `Operation Complete` to appear.
6. Quit the FTP session.
7. If the problem to diagnose seems specific to user-interface behavior, repeat [step 3](#) through [step 6](#) on the partner controller to collect its unique MC log data.

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 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" \(page 31\)](#).

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

To transfer log data to a log-collection system

1. In the SMC, prepare to use FTP:
 - a. Determine the network-port IP addresses of the system's controllers. See ["Changing network interface settings" \(page 54\)](#).

- b. Verify that the system's FTP service is enabled. See [“Changing system services settings”](#) (page 53).
 - c. Verify that the user you will log in as has permission to use the FTP interface. See [“To modify a user”](#) (page 46).
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:
`ftp controller-network-address`
For example:
`ftp 10.1.0.9`
4. Log in as a user that has permission to use the FTP 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`
Wait for the message `Operation Complete` to appear.
6. Quit the FTP session.

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

Downloading historical disk-performance statistics

You can access the storage system's FTP interface and use the `get perf` command to download historical disk-performance statistics for all disks in the storage system. This command downloads the data in CSV format to a file, for import into a spreadsheet or other third-party application.

The number of data samples downloaded is fixed at 100 to limit the size of the data file to be generated and transferred. The default is to retrieve all the available data (up to 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 client. A GUI-based FTP client might not work.

To retrieve historical disk-performance statistics

1. In the SMC, prepare to use FTP:
 - a. Determine the network-port IP addresses of the system's controllers. See [“Changing network interface settings”](#) (page 54).

- b. Verify that the system's FTP service is enabled. See [“Changing system services settings” \(page 53\)](#).
 - c. Verify that the user you will log in as has permission to use the FTP interface. See [“To modify a user” \(page 46\)](#).
2. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the destination directory for the log file.
 3. Enter:


```
ftp controller-network-address
```

 For example:


```
ftp 10.1.0.9
```
 4. Log in as a user that has permission to use the FTP 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
```

 Wait for the message `Operation Complete` to appear.
 6. Quit the FTP session.

Updating firmware

You can update the versions of firmware in controller modules, expansion modules (in drive enclosures), and disks.

 **TIP:** To ensure success of an online update, select a period of low I/O activity. This helps the update complete as quickly as possible and avoids disruptions to host and applications due to timeouts. Attempting to update a storage system that is processing a large, I/O-intensive batch job will likely cause hosts to lose connectivity with the storage system.

IMPORTANT:

- 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 116](#)).
-

Updating controller-module firmware

In a dual-controller system, both controllers should run the same firmware version. Storage systems in a replication set should run the same or compatible firmware versions. You can update the firmware in each controller module by loading a firmware file obtained from the Lenovo web download site at support.lenovo.com.


If you have a dual-controller system and the Partner Firmware Update (PFU) option is enabled (the default), when you update one controller the system automatically updates the partner controller. If PFU is disabled, after updating firmware on one controller you must log into the partner controller's IP address and perform this firmware update on that controller also.

For best results, ensure the storage system is in a healthy state before starting firmware update.

NOTE: For information about supported releases for firmware update, see the product's Release Notes.

To update controller module firmware

1. Obtain the appropriate firmware file and download it to your computer or network.
2. In the SMC, prepare to use FTP:
 - a. Determine the network-port IP addresses of the system's controllers.
 - b. Verify that the system's FTP service is enabled.
 - c. Verify that the user you will log in as has permission to use the FTP interface.
3. If the storage system has a single controller, stop I/O to disk groups before starting the firmware update.
4. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the directory containing the firmware file to load.
5. Enter:
`ftp controller-network-address`
For example:
`ftp 10.1.0.9`
6. Log in as an FTP user.
7. Enter:
`put firmware-file flash`
For example:
`put T230R01-01.bin 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 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 a drive enclosure.

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

If the Storage Controller cannot be updated, the update operation is canceled. If the FTP prompt does not return, quit the FTP session and log in again. Verify that you specified the correct firmware file and repeat the update. If this problem persists, contact technical support.

When firmware update on the local controller is complete, the message `Operation Complete` is printed, the FTP session returns to the `ftp>` prompt, and the FTP session to the local MC is closed.

If PFU is enabled, allow an additional 10–20 minutes for the partner controller to be updated.

8. Quit the FTP session.
9. Clear your web browser's cache, then sign in to the SMC. If PFU is running on the controller you sign in to, a dialog box shows PFU progress and prevents you from performing other tasks until PFU is complete.

NOTE: If PFU is enabled for the system through the `partner-firmware-upgrade` parameter of the `set advanced-settings` CLI command, after firmware update has completed on both controllers, check the system health. After firmware update has completed on both controllers, if the system health is Degraded and the health reason indicates that the firmware version is incorrect, verify that you specified the correct firmware file and repeat the update. If this problem persists, contact technical support.

Updating disk firmware

You can update disk firmware by loading a firmware file obtained from your reseller. A dual-ported disk can be updated from either controller.

NOTE: Disks of the same model in the storage system must have the same firmware revision.

You can specify to update all disks or only specific disks. If you specify to update all disks and the system contains more than one type of disk, the update will be attempted on all disks in the system. The update will only succeed for disks whose type matches the file, and will fail for disks of other types.

To prepare for update

1. Obtain the appropriate firmware file and download it to your computer or network.
2. Check the disk manufacturer's documentation to determine whether disks must be power cycled after firmware update.
3. If you want to update all disks of the type that the firmware applies to, continue with the next step. Otherwise, in the SMC, for each disk to update:
 - a. Determine the enclosure number and slot number of the disk.
 - b. If the disk is associated with a disk group and is single ported, determine which controller owns the disk group.
4. In the SMC, prepare to use FTP:
 - a. Determine the network-port IP addresses of the system's controllers.
 - b. Verify that the system's FTP service is enabled.
 - c. Verify that the user you will log in as has permission to use the FTP interface.
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. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the directory containing the firmware file to load.

2. Enter:
`ftp controller-network-address`
For example:
`ftp 10.1.0.9`
3. Log in as an FTP 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. Wait for a message that the update has succeeded.

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 session.
7. If the updated disks must be power cycled:
 - a. Shut down both controllers by using the SMC.
 - b. Power cycle all enclosures as described in your product's Hardware Installation and Maintenance Guide.
8. Verify that each disk has the correct firmware revision.

Installing a license file

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

`ftp controller-network-address`

For example:

`ftp 10.1.0.9`

4. Log in as an FTP user.

5. Enter:

`put license-file license`

For example:

`put certificate.txt license`

A message confirms whether installation succeeded or failed. If installation succeeds, licensing changes take effect immediately.

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 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 SMC, prepare to use FTP:
 - a. Determine the network-port IP addresses of the system's controllers. See [“Changing network interface settings” \(page 54\)](#).
 - b. Verify that the system's FTP service is enabled. See [“Changing system services settings” \(page 53\)](#).
 - c. Verify that the user you will log in as has permission to use the FTP interface. See [“To modify a user” \(page 46\)](#).
2. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the directory that contains the certificate files.
3. Enter:

```
ftp controller-network-address
```

For example:

```
ftp 10.1.0.9
```
4. Log in as a user that has permission to use the FTP 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 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 SMC, prepare to use FTP:
 - a. Determine the network-port IP addresses of the system's controllers. See [“Changing network interface settings” \(page 54\)](#).
 - b. Verify that the system's FTP service is enabled. See [“Changing system services settings” \(page 53\)](#).
 - c. Verify that the user you will log in as has permission to use the FTP interface. See [“To modify a user” \(page 46\)](#).
2. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the destination directory for the log file.

3. Enter:
`ftp controller-network-address`
For example:
`ftp 10.1.0.9`
4. Log in as a user that has permission to use the FTP interface.
5. Enter:
`get logs:heatmap filename.csv`
where: *filename* is the file that will contain the data.
For example:
`get logs:heatmap IO_density.csv`
Wait for the message `Operation Complete` to appear.
6. Quit the FTP session.

Using SLP

Lenovo ThinkSystem 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 22](#) and populate the configuration attributes shown in [Table 23](#).

Table 22 Interfaces advertised by SLP

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

Table 23 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
<code>x-system-name</code>	<code>system-name</code>
<code>x-system-contact</code>	<code>system-contact</code>
<code>x-system-location</code>	<code>system-location</code>
<code>x-system-information</code>	<code>system-information</code>

Table 23 SLP attributes shown for a storage system (continued)

SLP attribute	Corresponding property shown by the CLI <code>show system detail</code> command in XML API mode
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 SMC, as described in [“Changing system services settings” \(page 53\)](#), or by using the CLI `set protocols` command as described in the CLI Reference Guide.

If the SLP service is enabled, you can test it by using an open source tool, such as `slptool` from openSLP.org.

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” \(page 31\)](#).

Over time, a log-collection system can receive many log files from one or more storage systems. The administrator organizes and stores these log files on the log-collection system. Then, if a storage system experiences a problem that needs analysis, that system's current log data can be collected and combined with the stored historical log data to provide a long-term view of the system's operation for analysis.

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

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

Each log-file type also contains system-configuration information.

How log files are transferred and identified

Log files can be transferred to the log-collection system in two ways, depending on whether the managed logs feature is configured to operate in *push mode* or *pull mode*:

- In push mode, when log data has accumulated to a significant size, the storage system sends notification events with attached log files through email to the log-collection system. The notification specifies the storage-system name, location, contact, and IP address, and contains a single log segment in a compressed zip file. The log segment will be uniquely named to indicate the log-file type, the date/time of creation, and the storage system. This information will also be in the email subject line. The file name format is `logtype_yyyy_mm_dd_hh_mm_ss.zip`.
- In pull mode, when log data has accumulated to a significant size, the system sends notification events via email or SNMP traps, to the log-collection system. The notification will specify the storage-system name, location, contact, and IP address and the log-file type (region) that needs to be transferred. The storage system's FTP interface can be used to transfer the appropriate logs to the log-collection system, as described in [“Transferring log data to a log-collection system” \(page 133\)](#).

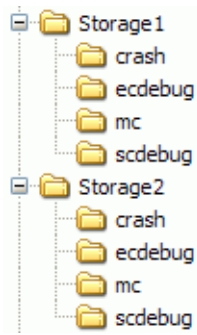
Log-file details

- 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 interface to get the log and save it into the `scdebug` subdirectory of the Storage2 directory.

C Best practices

This appendix describes best practices for configuring and provisioning a storage system.

Pool setup

In a storage system with two controller modules, try to balance the workload of the controllers. Each controller can own one virtual pool. Having the same number of disk groups and volumes in each pool will help balance the workload, increasing performance.

RAID selection

A pool is created by adding disk groups to it. Disk groups are based on RAID technology. The following table describes the characteristics and use cases of each RAID level.

Table 24 RAID level characteristics and use cases

RAID level	Protection	Performance	Capacity	Application use cases	Suggested disk speed
RAID 1/ RAID 10	Protects against up to one disk failure per mirror set	Great random I/O performance	Poor: 50% fault tolerance capacity loss	Databases, OLTP, Exchange Server, booting from SAN	10K, 15K, 7K
RAID 5	Protects against up to one disk failure per RAID set	Good sequential I/O performance, moderate random I/O performance	Great: One-disk fault tolerance capacity loss	Big data, media and entertainment (ingest, broadcast, and post production)	10K, 15K, lower capacity 7K
RAID 6	Protects against up to two disk failures per RAID set	Moderate sequential I/O performance, poor random I/O performance	Moderate: Two-disk fault tolerance capacity loss	Archive, parallel distributed file system	High capacity 7K

Disk count per RAID level

The controller breaks virtual volumes into 4-MB pages, which are referenced paged tables in memory. The 4-MB page is a fixed unit of allocation. Therefore, 4-MB units of data are pushed to a disk group. A write performance penalty is introduced in RAID-5 or RAID-6 disk groups when the stripe size of the disk group isn't a multiple of the 4-MB page.

- Example 1: Consider a RAID-5 disk group with five disks. The equivalent of four disks provide usable capacity, and the equivalent of one disk is used for parity (parity is distributed among disks). The four disks providing usable capacity are the data disks and the one disk providing parity is the parity disk. In reality, the parity is distributed among all the disks, but conceiving of it in this way helps with the example.

Note that the number of data disks is a power of two (2, 4, and 8). The controller will use a 512-KB stripe unit size when the data disks are a power of two. This results in a 4-MB page being evenly distributed across two stripes. This is ideal for performance.

- Example 2: Consider a RAID-5 disk group with six disks. The equivalent of five disks now provide usable capacity. Assume the controller again uses a stripe unit of 512-KB. When a 4-MB page is pushed to the disk group, one stripe will contain a full page, but the controller must read old data and old parity from two of the disks in combination with the new data in order to calculate new parity. This is known as a read-modify-write, and it's a performance killer with sequential workloads. In essence, every page push to a disk group would result in a read-modify-write.

To mitigate this issue, the controllers use a stripe unit of 64-KB when a RAID-5 or RAID-6 disk group isn't created with a power-of-two data disks. This results in many more full-stripe writes, but at the cost of many more I/O transactions per disk to push the same 4-MB page.

The following table shows recommended disk counts for RAID-6 and RAID-5 disk groups. Each entry specifies the total number of disks and the equivalent numbers of data and parity disks in the disk group. Note that parity is actually distributed among all the disks.

Table 25 Recommended disk group sizes

RAID level	Total disks	Data disks (equivalent)	Parity disks (equivalent)
RAID 6	4	2	2
	6	4	2
	10	8	2
RAID 5	3	2	1
	5	4	1
	9	6	1

To ensure best performance with sequential workloads and RAID-5 and RAID-6 disk groups, use a power-of-two data disks.

Disk groups in a pool

For better efficiency and performance, use similar disk groups in a pool.

- Disk count balance: For example, with 20 disks, it is better to have two 8+2 RAID-6 disk groups than one 10+2 RAID-6 disk group and one 6+2 RAID-6 disk group.
- RAID balance: It is better to have two RAID-5 disk groups than one RAID-5 disk group and one RAID-6 disk group.
- In terms of the write rate, due to wide striping, tiers and pools are as slow as their slowest disk groups.
- All disks in a tier should be the same type. For example, use all 10K disks or all 15K disks in the Standard tier.

Create more small disk groups instead of fewer large disk groups.

- Each disk group has a write queue depth limit of 100. This means that in write-intensive applications this architecture will sustain bigger queue depths within latency requirements.
- Using smaller disk groups will cost more raw capacity. For less performance-sensitive applications, such as archiving, bigger disk groups are desirable.

Tier setup

In general, it is best to have two tiers instead of three tiers. The highest tier will nearly fill before using the lowest tier. The highest tier must be 95% full before the controller will evict cold pages to a lower tier to make room for incoming writes.

Typically, you should use tiers with SSDs and 10K/15K disks, or tiers with SSDs and 7K disks. An exception may be if you need to use both SSDs and faster spinning disks to hit a combination of price for performance, but you cannot hit your capacity needs without the 7K disks; this should be rare.

VMware missing LUN response

For VMware environments, the missing LUN response setting should be set to *Illegal Request*. You can do this in the CLI by running the following command:

```
set advanced-settings missing-lun-response illegal
```

Multipath configuration

Lenovo ThinkSystem systems comply with the SCSI-3 standard for Asymmetrical Logical Unit Access (ALUA). ALUA compliant storage systems will provide optimal and non-optimal path information to the host during device discovery, but the operating system must be directed to use ALUA. You can use the following procedures to direct Windows and Linux systems to use ALUA.

To direct a Windows 2012 system to use ALUA

1. Verify that MPIO is enabled on the host by starting Server Manager.
 - a. Select **Local Server**.
 - b. Select **Add Roles and Features**.
 - c. Select **Role-based or Feature Based Installation**.
 - d. Select the server from the pool.
 - e. Click Next to navigate to the feature selection screen.
 - f. Verify that Multipath IO is installed. If it is not installed, click the box and then select **Install**.
2. When MPIO is installed, navigate to Server Manager, then Tools.
3. Select **MPIO** from the menu.
4. Select **Discover Multi-Paths**.
5. Your storage system model should be listed. Select the device and then **Add**.
6. When prompted, reboot the system. When the reboot is complete, the system is ready to use.

To direct a Linux system to use ALUA

1. Ensure that the multipath daemon is installed and set to start at run-time. Linux command:
`chkconfig multipathd on`
2. Ensure the correct entries exist in the `/etc/multipath.conf` file on each OSS/MDS host. Create a separate device entry for the Lenovo ThinkSystem system. The following table specifies four attributes that should be set. To obtain the exact vendor and product ID values, run the Linux command: `multipath -v3`

Attribute	Value
prio	alua
failback	immediate
vendor	Lenovo
product	product-ID

3. Instruct the multipath daemon to reload the `multipath.conf` file or reboot the server. Linux command:
`service multipathd reload`
4. Determine if the multipath daemon used ALUA to obtain the optimal/non-optimal paths. Linux command:
`multipath -v3 | grep alua`
You should see output stating that ALUA was used to configure the path priorities. For example:
`Oct 01 14:28:43 | sdb: prio = alua (controller setting) Oct 01 14:28:43 | sdb: alua prio = 130`

Physical port selection

In a system configured to use either all FC or all iSCSI ports, use the ports in the following order:

1. A0,B0
2. A2,B2
3. A1,B1
4. A3,B3

The reason for doing so is that each pair of ports (A0,A1 or A2,A3) are connected to a dedicated CNC chip. If you are not using all four ports on a controller, it is best to use one port from each pair (A0,A2) to ensure better I/O balance on the front end.

Boot from SAN

- Use RAID 10 to boot from SAN environments.
- Use faster disks (10K and 15K) to deliver the required performance.
- Data LUNs should be separate from boot LUNs.
- Data LUNs can be used for heavy I/O but boot LUNs should only be used for operating systems.
- As a rule of thumb, use 12 disks in a RAID-10 configuration to run 7 operating systems (7 BFS LUNs + 7 data LUNs).

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.
ATS	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.
CNC	Converged Network Controller. A controller module whose host ports can be set to operate in FC or iSCSI mode, using qualified SFP and cable options. Changing the host-port mode is also known as changing the ports' personality.
compatible disk	A disk that can be used to replace a failed member disk of a disk group because it both has enough capacity 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.
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 type and provides storage capacity for a pool. <i>See also</i> virtual disk group.
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, MC, SC.
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	See EC.
expansion enclosure	An enclosure that contains one or two expansion modules. Expansion enclosures can be connected to a controller enclosure to provide additional storage capacity. See <i>also</i> EBOD, JBOD.
expansion module	A FRU that contains the following subsystems and devices: a SAS expander and EC 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. See <i>also</i> default mapping, masking.
failback	See recovery.
failover	In an active-active configuration, failover is the act of temporarily transferring ownership of controller resources from an offline controller to its partner controller, which remains operational. The resources include pools, volumes, cache data, host ID information, and LUNs and WWNs. See 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. See <i>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. See CRU.
full disk encryption	See 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. See <i>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 or switch.
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	See IOM.
IOM	Input/output 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." See 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. <i>See also</i> FDE, passphrase.
loop	<i>See</i> 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	<i>See</i> 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. <i>See also</i> default mapping, explicit mapping, masking.
masking	A volume-mapping setting that specifies no access to that volume by hosts. <i>See 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 SMC, and computer-computer interfaces, such as SNMP, and interacts with the Storage Controller. <i>See also</i> EC, SC.
metadata	Data in the first sectors of a disk drive that stores all 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. <i>See 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	<i>See</i> 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	<i>See</i> PFU.
passphrase	A user-created password that allows users to manage lock keys in an FDE-capable system. <i>See 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. <i>See 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 iSCSI links. The capability to replicate volumes is a licensed feature.
replication set	For virtual 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.
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. A type of disk drive.
SHA	Secure Hash Algorithm 1 authentication.
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.
SMC	Storage Management Console. 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. 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	<i>See</i> SC.
Storage Management Console	<i>See</i> SMC.
storage system	A controller enclosure with at least one connected drive 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.

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 Asymmetric Logical Unit Access (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.
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 and SMC 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 type. The number of disks that a virtual disk group can contain is determined by its RAID level. A virtual disk group can use RAID 1, 5, 6, or 10. A virtual disk group can be added to a new or existing virtual pool. See also 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 Console. The primary interface for managing the system. A user can enable the use of HTTP, HTTPS for increased security, or both.
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 13

A

activity progress interface 61

all-flash array

about 21

allocated space

virtual storage 119

ALUA 27

archive tier 25

asterisk (*) in option name 13

audience 10

Automated Tiered Storage (ATS)

about 25

advantages 26

frequently accessed data 25

infrequently accessed data 25

B

banner

overview 114

base for size representations 16

best practices 144

bytes versus characters 16

C

cache

configuring volume settings 85

Capacity block

physical and logical storage identification 35

capacity information 35

viewing 118

Capacity Utilization panel 35, 119

certificate

using FTP to install a security 139

CHAP

adding or modifying a CHAP record 71

configuring 56

configuring for iSCSI hosts 70

configuring through Configuration Wizard 42

deleting a CHAP record 71

overview 26

setting up for use with a peer connection 103

using in a system with a peer connection 71

using with replication 103

characters versus bytes 16

color codes for storage space 15

configuration

browser 12

first-time 12

Configuration Wizard, using 36

controllers

restarting or shutting down 65

using FTP to update firmware 135

using the SMC to update firmware 59

conventions

document 11

Critical & Error Event Information panel 117

D

date and time

changing through Configuration Wizard 37

configuring 115

debug data

saving to a file 116

debug logs

downloading 132

default mapping

about 27

advantages and disadvantages 92

DHCP

configuring 54

configuring with Configuration Wizard 38

disaster recovery

accessing data from backup system 100

accessing data with intact replication set 100

procedures 100

using replication in 99

disk channels

rescanning 57

disk groups

about 18, 22

adding 74

listed information 72

modifying 76

options 75

read-cache 19, 75

removing 77

scrubbing 79

using SSDs 75

virtual 18, 75

disk metadata

clearing 57

disk sector format

identifying 74

disks

identifying sector format 74

using FTP to retrieve performance statistics 134

using FTP to update firmware 137

using the SMC to update firmware 60

- document
 - conventions 11
 - prerequisite knowledge 10
 - related documentation 10
- DWD
 - SSD endurance indicator 22

E

- empty allocated pages
 - replication 99
- enclosure
 - front view 50
 - rear view 50
 - table view 51
 - viewing information about 50
- enclosure properties 51
- event log
 - viewing 118
- event notification
 - changing settings 46
 - configuring email settings 47
 - configuring SNMP settings 46
 - configuring with Configuration Wizard 40
 - testing settings 48
- event severity icon 118
- explicit mapping
 - about 27, 92

F

- FDE
 - about 32
 - changing settings 62
 - clearing lock keys 63
 - repurposing disks 64
 - repurposing system 64
 - securing the system 63
 - setting FDE import lock key IDs 65
 - setting the passphrase 62
- firmware
 - about updating 30
 - updating through FTP 135
 - updating through the SMC 58
 - updating, best practices 58
 - using FTP to update controller module firmware 135
 - using FTP to update disk drive firmware 137
 - using the SMC to update controller module firmware 59
 - using the SMC to update disk firmware 60
 - using the SMC to update the expansion module 60
- firmware update, monitoring progress 61
- footer
 - overview 114
- foreign virtual disk group
 - resolving a resulting pool conflict 36

- FTP
 - downloading system logs 132
 - retrieving disk-performance statistics 134
 - updating controller module firmware 135
 - updating disk drive firmware 137
 - updating firmware 135
 - using to download system heat map data 139
 - using to install a security certificate 139
 - using with the log-management feature 133

- Full Disk Encryption

 - See FDE

G

- global spares 52
 - adding and removing 52
- grouping
 - maximum number of hosts 26
 - maximum number of initiators 26

H

- heat map data
 - using FTP to download 139
- help
 - using online 14
- historical performance statistics
 - exporting 112
 - graphs 110
 - resetting 113
 - updating 112
- Home topic
 - host information 34
 - IOPS port information 35
 - port data throughput information 35
 - port information 34
 - spares information 36
 - storage capacity information 35
 - system health information 36
 - viewing system status 34
- host
 - adding initiators to 68
 - changing name 69
 - definition 34
 - removing initiators 69
 - viewing information about 67
- host group
 - adding hosts 69
 - definition 34
 - removing hosts 70
 - renaming 70
- host groups
 - about 26
 - mapping 92
 - removing 70
 - viewing 67

- host I/O information
 - viewing 119
- host ports
 - configuring 55
 - configuring with Configuration Wizard 41
 - resetting 57
- hosts
 - about 26
 - adding to host group 69
 - basic information 67
 - list of 67
 - mapping 92
 - maximum number in a host group 26
 - removing 69
 - removing from host group 70

I

- icons
 - event severity 118
 - SMC communication status 115

- initiator
 - definition 34
 - deleting 68
 - manual creation 68
 - modifying 68
 - nickname 26

- initiators
 - about 26
 - adding to a host 68
 - mapping 92, 93
 - maximum number in a host 26
 - removing from a host 69
 - viewing 67

- iSCSI host security 26

- iSCSI IP version
 - configuring 56
 - configuring through Configuration Wizard 42

- iSNS
 - configuring 56
 - configuring through Configuration Wizard 42

J

- jumbo frames
 - configuring 56
 - configuring through Configuration Wizard 42

L

- large pools
 - snapshot limits 88
- leftover disk 58

- licensed features
 - about 49
 - creating a temporary license 49
 - installing a permanent license 49
 - snapshot limit 28
 - using FTP to install license file 138
 - viewing status of 49
 - volume copy 29

- lock key 32
- log data
 - saving to a file 116
- log management
 - about 31
 - using FTP 133
- log-collection system
 - administering 142
- logs
 - downloading debug 132
- LUNs
 - about 27

M

- managed logs
 - about 31
 - administering a log-collection system 142
 - pull mode 31
 - push mode 31
- management interface services
 - configuring with Configuration Wizard 39
- mapping volumes
 - See volume mapping
- metadata
 - clearing disk 57
- MIB
 - See SNMP

N

- network ports
 - configuring 54
 - configuring with Configuration Wizard 38
- nickname
 - initiator 26
- notification history
 - viewing 120
- NTP
 - configuring 115

O

- overcommitment setting
 - enabling 78
- overcommitting physical storage
 - about 25

P

- passphrase 32
- passwords
 - See users
- peer connection
 - CHAP setup 103
- peer connections
 - creating 102
 - deleting 104
 - modifying 104
 - table 101
- performance monitoring
 - See storage system component performance
- performance statistics
 - about 30
 - historical performance graphs 110
 - resetting 113
 - viewing 110
- performance tier 25
- pools
 - about 22
 - attributes 72
 - changing settings 78
 - large pools snapshot limits 88
 - list of 72
 - viewing information about 72
 - virtual 23
- ports
 - attributes and status 34
 - data throughput 35
 - IOPS information 35
- prerequisite knowledge 10
- provisioning
 - first-time 12

Q

- quick rebuild
 - about 29

R

- read cache
 - about 22
 - advantages 22
 - cache utilization graph 36
- read-ahead caching
 - Adaptive option 25
 - Disabled option 25
 - optimizing 24
 - Stripe option 25
- read-cache disk groups
 - about 19, 23
- related documentation 10

- replication
 - about 96
 - creating a virtual pool for 99
 - of empty allocated pages 99
 - prerequisites 96
 - process 96
 - setting up snapshot space management for 99
 - using CHAP with 103
 - using in disaster recovery 99
- replication process 96
 - initial replication 97
 - internal snapshot space 98
 - subsequent replications 98
- replication sets
 - creating from the Replications topic 105
 - creating from the Volumes topic 90
 - deleting 106
 - modifying 106
 - primary volumes and volume groups 105
 - secondary volumes and volume groups 105
- replications
 - aborting 108
 - initiating from the Replications topic 106
 - initiating from the Volumes topic 91
 - manage replication schedule 108
 - Peer Connections table 101
 - Replication Sets table 101
 - scheduling 107
 - viewing 101
- repurposing
 - disks 64
 - secured disks and systems 32
 - system 64
- rescan disk channels 57
- reserved space 119
- resetting host ports 57
- restarting controllers 65
- rolling back volume data
 - about 87
 - virtual volumes and snapshots 87

S

- schedules
 - deleting 48
 - manage replication schedules 108
 - managing 48
 - modifying 48
 - see information about 83
- scheduling
 - replications 107
 - snapshot 88
 - snapshot reset 89
- sector format
 - identifying 74

- security certificate
 - using FTP to install 139
- shutting down controllers 65
- sign out, auto
 - viewing remaining time 13
- signing in to the SMC 17
- single-controller system data-protection tips 33
- size representations
 - about 16
- SLP
 - advertised interfaces 140
 - attributes 140
 - overview 140
- SMC
 - about 12
 - communication status icon 115
 - signing in 17
- snapshot
 - basic information 81
 - creating 88
 - resetting to current data in source volume 89
- snapshot space management
 - in the context of replication 99
- snapshots
 - about 28
 - about reset snapshot 28
 - about rollback 28
 - automatic deletion 28
 - deleting 88
 - list of 81
 - list of child snapshots 81
 - mapping 92
 - resetting 28
 - rollback feature 28
 - setting snapshot pool space 28
- snapshots, virtual
 - creating 88
 - levels 28
 - parent-child relationships 28
 - rollback feature 28
 - snapshot hierarchy 28
- SNMP
 - configuring traps 130
 - enterprise trap MIB 130
 - enterprise traps 121
 - external details for connUnitPortTable 129
 - external details for connUnitRevsTable 127
 - external details for connUnitSensorTable 128
 - FA MIB 2.2 behavior 122
 - FA MIB 2.2 objects, descriptions, and values 122
 - management 130
 - MIB-II behavior 121
 - overview 121
 - setting event notification 130
- sorting a table 14
- Spares
 - about 22
 - rules for 22
 - types of 22
- spares
 - Home topic information 36
- SSD read cache
 - about 22
- SSDs
 - about 20
 - cost/benefit analysis 20
 - data retention 22
 - DWD 22
 - endurance indicated by DWD 22
 - gauging percentage of life remaining 21
 - internal disk management 21
 - maintenance 20
 - overprovisioning 21
 - rules for using 20
 - SSD Life Left disk property 21
 - TRIM and UNMAP commands 21
 - using in a disk group 75
 - wear leveling 21
- standard tier 25
- storage blocks 35
 - logical storage information 35
 - read cache 36
 - virtual storage 35
- Storage Management Console
 - about 12
- storage system
 - See system
- storage system component performance
 - about monitoring historical data 30
- system
 - data-protection tips for a single-controller 33
- system activity
 - viewing 119
- system components
 - properties 51
- system health 36
 - viewing 116
- System Health panel 116
- system information
 - configuring 43
 - configuring with Configuration Wizard 40
 - menu options 114
 - viewing 114
- System Information panel 114
- system service settings
 - changing 53
- system status
 - viewing 34

T

- table sorting 14
- tables
 - tips for using 14
- task schedule
 - See schedule
- thin provisioning
 - about 25
 - overcommit storage 25
- tiers
 - archive 25
 - performance 25
 - standard 25
 - viewing I/O information 119
- time and date
 - configuring 115
- troubleshooting resources 118

U

- ULP 27
- unallocated space
 - virtual storage 119
- units for size representations 16
- User Information panel 116
- user interface
 - main areas 12
- user panel
 - changing user settings 116
- users
 - adding 45
 - changing default passwords with Configuration Wizard 38
 - deleting 45
 - modifying 45

V

- virtual disk groups
 - about 18
 - number allowed per pool 18
 - removal requirements 77
 - requirements 18
- virtual pools
 - about 22
 - about removing 23
 - volume allocation 23
- virtual snapshots
 - about 28
 - creation process 28
- virtual storage
 - advantages 17
 - page definition 17
 - quick rebuild 29
 - reconstruction using quick-rebuild features 29
- virtual volume
 - creating 84

- virtual volumes
 - about adding to virtual pools 23
- volume
 - basic information 81
 - changing name 85
 - configuring cache settings 85
 - expanding 85
 - modifying 85
 - rolling back data 87
 - viewing information about 81
- volume cache options
 - about 23
- volume copy
 - about 29
- volume creation
 - default mapping 27
- volume group
 - adding volumes 85
 - removing 86
 - removing group and volumes 86
 - removing volumes from 86
 - renaming 86
- volume groups
 - about 23
 - mapping 92
 - maximum number of volumes 23
 - requirements 23
- volume mapping
 - about 27
 - editing 94
 - procedure 93
 - Related Maps table 67
 - unmapping 94
 - viewing details 95
 - viewing information about 67, 92
- volume tier affinity
 - about 26
 - settings 26
- volumes
 - about 23
 - adding to volume group 85
 - creating a virtual volume 84
 - deleting 88
 - list of 81
 - mapping 92
 - removing from a volume group 86
 - virtual 23
- Volumes topic
 - Maps table 82
 - Replication Sets table 82
 - Schedules table 83
 - Snapshots table 81
 - volumes table 81

W

- web-browser buttons to avoid 13
- web-browser setup 12
- write-back caching 24
- write-through caching 24