



Dot Hill AssuredSAN 4004 Series RAIDar User Guide

Abstract

This guide is for use by storage administrators to manage a Dot Hill AssuredSAN 4004 Series storage system by using its web interface, RAIDar Storage Management Utility.

Copyright © 2013 Dot Hill Systems Corp. All rights reserved. Dot Hill Systems Corp., Dot Hill, the Dot Hill logo, AssuredSAN, AssuredSnap, AssuredCopy, AssuredRemote, EcoStor, SimulCache, R/Evolution, and the R/Evolution logo are trademarks of Dot Hill Systems Corp. All other trademarks and registered trademarks are proprietary to their respective owners.

The material in this document is for information only and is subject to change without notice. While reasonable efforts have been made in the preparation of this document to assure its accuracy, changes in the product design can be made without reservation and without notification to its users.



Adobe PostScript

Contents

About this guide	13
Intended audience	13
Prerequisites	13
Related documentation	13
Document conventions and symbols	14
1 Getting started	15
Configuring and provisioning a new storage system	15
Browser setup	15
Signing in	15
Tips for signing in and signing out	16
Tips for using the main window	16
Tips for using the help window	17
System concepts	17
About user accounts	17
About vdisks	18
About spares	19
About volumes	20
About hosts	21
About volume mapping	22
About volume cache options	22
About managing remote systems	24
About the Snapshot feature	24
About the Volume Copy feature	26
About the AssuredRemote replication feature	28
About the VDS and VSS hardware providers	28
About the Storage Replication Adapter (SRA)	28
About RAID levels	28
About size representations	30
About the system date and time	30
About storage-space color codes	31
About Configuration View icons	31
About disk failure and vdisk reconstruction	31
About data protection in a single-controller storage system	32
About managed logs	32
About performance monitoring	34
About firmware update	34
2 Configuring the system	37
Using the Configuration Wizard	37
Step 1: Starting the wizard	37
Step 2: Changing default passwords	37
Step 3: Configuring network ports	37
Step 4: Enabling system-management services	38
Step 5: Setting system information	39
Step 6: Configuring event notification	39
Step 7: Configuring host ports	40
Step 8: Confirming configuration changes	41
Installing a license	41
Configuring system services	42
Changing management interface settings	42
Configuring email notification	43
Configuring SNMP notification	44

Configuring user accounts	44
Adding users.	44
Modifying users.	45
Removing users	47
Configuring system settings	47
Changing the system date and time	47
Changing host interface settings	47
Changing network interface settings.	49
Setting system information	50
Configuring advanced settings	50
Changing disk settings	50
Changing system cache settings	52
Configuring partner firmware update	54
Configuring system utilities	54
Configuring remote systems	55
Adding a remote system	55
Deleting remote systems	55
Configuring a vdisk	56
Managing dedicated spares	56
Changing a vdisk's name	56
Changing a vdisk's owner	57
Configuring drive spin down for a vdisk	57
Configuring a volume	58
Changing a volume's name	58
Changing a volume's cache settings	58
Configuring a snapshot	58
Changing a snapshot's name	58
Configuring a snap pool	58
Changing a snap pool's name	58
3 Provisioning the system	59
Using the Provisioning Wizard.	59
Step 1: Starting the wizard	59
Step 2: Specifying the vdisk name and RAID level	59
Step 3: Selecting disks	60
Step 4: Defining volumes	60
Step 5: Setting the default mapping	60
Step 6: Confirming vdisk settings.	61
Creating a vdisk	61
Deleting vdisks	62
Managing global spares	62
Creating a volume set	63
Creating a volume	63
Deleting volumes	64
Changing default mapping for multiple volumes	64
Explicitly mapping multiple volumes	65
Changing a volume's default mapping	66
Changing a volume's explicit mappings	67
Unmapping volumes	68
Expanding a volume	68
Creating multiple snapshots	68
Creating a snapshot	69
Deleting snapshots	70
Resetting a snapshot	70
Creating a volume copy	71
Aborting a volume copy	73
Rolling back a volume	73
Creating a snap pool	74
Deleting snap pools	74

Adding a host	74
Removing hosts	75
Changing a host's name	75
Changing host mappings.	75
Configuring CHAP	76
Modifying a schedule	77
Deleting schedules	78
4 Using system tools	79
Updating firmware	79
Updating controller-module firmware	79
Updating expansion-module firmware	80
Updating disk firmware	81
Using the activity progress interface	81
Saving logs	83
Resetting a host port	83
Rescanning disk channels	84
Restoring system defaults	84
Clearing disk metadata	84
Restarting or shutting down controllers.	85
Restarting	85
Shutting down	86
Testing notifications.	86
Expanding a vdisk	86
Verifying a vdisk.	87
Scrubbing a vdisk.	87
Removing a vdisk from quarantine	88
Expanding a snap pool	89
Checking links to a remote system.	90
Checking local system links	90
Resetting or saving historical disk-performance statistics.	90
Resetting historical disk-performance statistics	90
Saving historical disk-performance statistics.	91
5 Viewing system status	93
Viewing information about the system	93
System properties	93
Enclosure properties.	94
Disk properties	94
Vdisk properties.	95
Volume properties	96
Snap-pool properties	96
Snapshot properties	96
Schedule properties	96
Configuration limits	97
Licensed features	97
Version properties	97
Viewing the system event log	98
Viewing information about all vdisks.	99
Viewing information about a vdisk	100
Vdisk properties.	100
Vdisk performance.	101
Disk properties	102
Volume properties	103
Snap-pool properties	103
Viewing information about a volume.	104
Volume properties	104
Mapping properties	105
Schedule properties	105

Viewing information about a snapshot	106
Snapshot properties	106
Mapping properties	107
Schedule properties	108
Viewing information about a snap pool.	108
Snap-pool properties	108
Volume properties	109
Snapshot properties	109
Viewing information about all hosts	109
Viewing information about a host.	110
Host properties	110
Mapping properties	110
Viewing information about an enclosure	110
Enclosure properties.	111
Disk properties	111
Disk performance	113
Power supply properties	114
Controller module properties.	114
Controller module: network port properties	115
Controller module: FC host port properties	115
Controller module: iSCSI host port properties	116
Controller module: SAS host port properties	117
Controller module: expansion port properties	117
Controller module: CompactFlash properties.	117
Drive enclosure: I/O module properties	118
I/O module: In port properties	118
I/O module: Out port properties	118
Viewing information about a remote system.	119
6 Using AssuredRemote to replicate volumes	121
About the AssuredRemote replication feature	121
Replication process overview	121
Replication actions.	123
Performing initial replication locally or remotely.	124
Criteria for selecting a vdisk to contain a secondary volume	124
Remote replication disaster recovery	125
Remote replication licensing	126
Related topics	126
Using the Replication Setup Wizard	127
Step 1: Starting the wizard	127
Step 2: Selecting the primary volume	127
Step 3: Selecting the replication mode	128
Step 4: Selecting the secondary volume	128
Step 5: Confirming replication settings	128
Replicating a volume.	129
Replicating a snapshot	130
Removing replication from a volume	131
Suspending replication	131
Resuming replication.	132
Aborting replication	132
Detaching a secondary volume	132
Stopping a vdisk	133
Starting a vdisk	133
Reattaching a secondary volume	134
Exporting a replication image to a snapshot	135
Changing the primary volume for a replication set	135

Viewing replication properties, addresses, and images for a volume.	136
Replication properties.	136
Replication addresses.	137
Replication images.	138
Viewing information about a remote primary or secondary volume.	138
Replication properties.	138
Replication addresses.	138
Replication image properties.	138
Viewing information about a replication image	138
Replication status properties	139
Primary volume snapshot properties	139
Secondary volume snapshot properties.	139
A SNMP reference.	141
Supported SNMP versions	141
Standard MIB-II behavior	141
Enterprise traps	141
FA MIB 2.2 SNMP behavior	142
External details for certain FA MIB 2.2 objects	147
External details for connUnitRevsTable	147
External details for connUnitSensorTable.	148
External details for connUnitPortTable.	149
Configuring SNMP event notification in RAIDar	149
SNMP management	149
Enterprise trap MIB	150
B Using FTP	153
Downloading system logs	153
Transferring log data to a log-collection system.	154
Downloading historical disk-performance statistics	155
Updating firmware	156
Updating controller-module firmware	156
Updating expansion-module firmware	157
Updating disk firmware	159
Installing a license file	160
Installing a security certificate.	160
C Using SMI-S.	163
Embedded SMI-S array provider.	163
SMI-S implementation.	164
SMI-S architecture	164
About the AssuredSAN 4004 Series SMI-S provider	165
SMI-S profiles.	166
Block Server Performance subprofile.	167
CIM	167
Supported CIM operations	167
CIM Alerts	167
Life cycle indications	168
SMI-S configuration.	169
Listening for managed-logs notifications.	169
Testing SMI-S	170
LUN Masking and Mapping operations.	170
Troubleshooting	170
D Administering a log-collection system	171
How log files are transferred and identified	171
Log-file details	171
Storing log files	172

Glossary	173
Index	181

Figures

1 Relationship between a master volume and its snapshots and snap pool	24
2 Rolling back a master volume	25
3 Creating a volume copy from a master volume or a snapshot	27
4 Intersite and intrasite replication sets	122
5 Actions that occur during a series of replications	123
6 Example of primary-volume failure	126

Tables

1	Document conventions	14
2	RAIDar communication status icons	16
3	Settings for default users	18
4	Example applications and RAID levels	28
5	RAID level comparison	29
6	Vdisk expansion by RAID level	29
7	Size representations in base 2 and base 10	30
8	Decimal (radix) point character by locale	30
9	Storage-space color codes	31
10	Configuration View icons	31
11	Activity progress properties and values	82
12	Available space required for a vdisk to be selectable to contain a secondary volume	125
13	FA MIB 2.2 objects, descriptions, and values	142
14	connUnitRevsTable index and description values	147
15	connUnitSensorTable index, name, type, and characteristic values	148
16	connUnitPortTable index and name values	149
17	Supported SMI-S profiles	166
18	CIM Alert indication events	167
19	Life cycle indications	168
20	Troubleshooting	170

About this guide

This guide provides information about managing a AssuredSAN™ 4004 Series storage system by using its web interface, RAIDar Storage Management Utility.

Intended audience

This guide is intended for storage administrators.

Prerequisites

Prerequisites for using this product include knowledge of:

- Network administration
- Storage system configuration
- Storage area network (SAN) management and direct attach storage (DAS)
- Fibre Channel, Serial Attached SCSI (SAS), Internet SCSI (iSCSI), and Ethernet protocols

Related documentation

For information about	See
Enhancements, known issues, and late-breaking information not included in product documentation	Release Notes
Overview of product shipkit contents and setup tasks	Getting Started*
Regulatory compliance and safety and disposal information	AssuredSAN Product Regulatory Compliance and Safety*
Installing and using optional host-based software components (CAPI Proxy, MPIO DSM, VDS Provider, VSS Provider, SES Driver, SRA)	AssuredSAN Installing Optional Software for Microsoft Windows® Server
Using a rackmount bracket kit to install an enclosure into a rack	AssuredSAN Rackmount Bracket Kit Installation* or AssuredSAN 2-Post Rackmount Bracket Kit Installation*
Product hardware setup and related troubleshooting	AssuredSAN 4004 Series Setup Guide
Obtaining and installing a license to use licensed features	AssuredSAN 4004 Series Obtaining and Installing a License
Using the web interface to configure and manage the product	AssuredSAN 4004 Series RAIDar User Guide
Using the command-line interface (CLI) to configure and manage the product	AssuredSAN 4004 Series CLI Reference Guide
Event codes and recommended actions	AssuredSAN Event Descriptions Reference Guide
Identifying and installing or replacing field-replaceable units (FRUs)	AssuredSAN 4004 Series FRU Installation and Replacement Guide

* Printed document included in product shipkit.

For additional information, see Dot Hill's Customer Resource Center web site: <http://crc.dothill.com>.

Document conventions and symbols

Table 1 Document conventions

Convention	Element
Blue text	Cross-reference links and e-mail addresses
Blue, underlined text	Web site addresses
Bold font	<ul style="list-style-type: none">• Key names• Text typed into a GUI element, such as into a box• GUI elements that are clicked or selected, such as menu and list items, buttons, and check boxes
<i>Italics font</i>	Text emphasis
Monospace font	<ul style="list-style-type: none">• File and directory names• System output• Code• Text typed at the command-line
<i>Monospace, italic font</i>	<ul style="list-style-type: none">• Code variables• Command-line variables
Monospace, bold font	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

RAIDar is a web-based application for configuring, monitoring, and managing the storage system.

Each controller module in the storage system contains a web server, which is accessed when you sign in to RAIDar. 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.

RAIDar is also referred to as the web-browser interface (WBI).

Configuring and provisioning a new storage system

To configure and provision a storage system for the first time:

1. Configure your web browser to access RAIDar and sign in, as described in [Browser setup](#) and [Signing in](#) below.
2. Set the system date and time, as described in ["Changing the system date and time"](#) (page 47).
3. Use the Configuration Wizard to configure other system settings, as described in ["Using the Configuration Wizard"](#) (page 37).
4. Use the Provisioning Wizard to create a virtual disk (*vdisk*) containing storage volumes, and optionally to map the volumes to hosts, as described in ["Using the Provisioning Wizard"](#) (page 59).
5. Use the Replication Setup Wizard to configure replication for a primary volume to a remote system, as described in ["Using the Replication Setup Wizard"](#) (page 127).
6. If you mapped volumes to hosts, verify the mappings by mounting/presenting the volumes from each host and performing simple read/write tests to the volumes.
7. Verify that controller modules and expansion modules have the latest firmware, as described in ["Viewing information about the system"](#) (page 93) and ["Updating firmware"](#) (page 79).

You can make additional configuration and provisioning changes and view system status, as described in later chapters of this guide.

Browser setup

- Supported browser versions: Mozilla Firefox 11 and newer; Microsoft Internet Explorer 9 and 10; Google Chrome 17 and newer; 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):
 - 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 RAIDar is configured to use HTTPS, ensure that Internet Explorer is set to use SSL 3.0 and either TLS 1.1 or TLS 1.2.

Signing in

To sign in:

1. In the web browser's address field, type the IP address of a controller network port and press **Enter**. The RAIDar 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 System Overview page is displayed; otherwise, a message indicates that the system is unavailable.




Tips for signing in and signing out

- Do not include a leading zero in an IP address. For example, enter 10.1.4.33 not 10.1.4.033.
- Multiple users can be signed in to each controller simultaneously.
- For each active RAIDar 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. Each instance of Internet Explorer can run a separate RAIDar session; however, all instances of Firefox share the same session.
- End a RAIDar session by clicking the Sign Out link near the top of the RAIDar window. Do not simply close the browser window.

Tips for using the main window




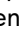
- The Configuration View panel displays logical and physical components of the storage system. To perform a task, select the component to act on and then either:
 - Right-click to display a context menu and select the task to perform. This is the method that help topics describe.
 - Click a task category in the main panel and select the task to perform.
- The System Status panel shows the system time and how many events of each severity have occurred in the system. To view event details, click a severity icon. For more information see ["Viewing the system event log" \(page 98\)](#).
- Many tables can be sorted by a specific column. To do so, click the column heading to sort low to high; click again to sort high to low. In tables that allow a task to be performed on multiple items, you can select up to 100 items or clear all selections by toggling the checkbox in the table's heading row.
- Do not use the browser's Back, Forward, Reload, or Refresh buttons. RAIDar has a single page whose content changes as you perform tasks and automatically updates to show current data.
- A red asterisk (*) identifies a required setting.
- The icon in the upper right corner of the main window shows the status of communication between RAIDar, the Management Controller (MC), and the Storage Controller (SC), as described in the following table.

Table 2 RAIDar communication status icons

Icon	Meaning
	RAIDar can communicate with the Management Controller, which can communicate with the Storage Controller.
	RAIDar <i>cannot</i> communicate with the Management Controller.
	RAIDar can communicate with the Management Controller, which <i>cannot</i> communicate with the Storage Controller.

- Below the communication status icon, a timer shows how long the session can be idle until you are automatically signed out. This timer resets after each action you perform. One minute before automatic sign-out you are prompted to continue using RAIDar.
- If a RAIDar session is active on a controller and the controller is power cycled or is forced offline by the partner controller or certain other events occur, the session might hang. RAIDar might say that it is "Connecting" but stop responding, or the page may become blank with the browser status "Done." After the controller comes back online, the session will not restart. To continue using RAIDar, close and reopen the browser and start a new RAIDar session.
- Colors that identify how storage space is used are described in ["About storage-space color codes" \(page 31\)](#).
- Icons shown in the Configuration View panel are described in ["About Configuration View icons" \(page 31\)](#).

Tips for using the help window

- To display help for a component in the Configuration View panel, right-click the component and select Help. To display help for the content in the main panel, click either Help in the menu bar or the help icon  in the upper right corner of the panel.
- 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 help window and you can open a new help 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.

System concepts

About user accounts

The system provides three default user accounts and allows a maximum of 12 user accounts to be configured. Any account can be modified or removed except you cannot remove the user you are signed in as.

The default user accounts are for general users that can access the WBI (RAIDar), CLI, FTP, or SMI-S interfaces. You can also create SNMPv3 user accounts that can access the Management Information Base (MIB) or receive trap notifications. SNMPv3 user accounts support SNMPv3 security features such as authentication and encryption. For information about configuring trap notifications, see ["Configuring SNMP notification" \(page 44\)](#). For information about the MIB, see ["SNMP reference" \(page 141\)](#).

General user accounts have these options:

- User Name.
- Password.
- User Roles. Either: Monitor, which lets the user view system settings; or Manage, which lets the user view and change system settings.
- User Type. Identifies the user's experience level: Standard, Advanced, or Diagnostic. This parameter does not affect access to the commands.
- WBI Access. Allows access to the web-based management interface.
- CLI Access. Allows access to the command-line management interface.
- FTP Access. Allows access to the FTP interface, which can be used instead of the WBI to install firmware updates and download logs.
- SMIS Access. Allows access to the Storage Management Initiative Specification (SMI-S) interface, used for management of the system through your network.
- Base Preference. The base for entry and display of storage-space sizes. In base 2, sizes are shown as powers of 2, using 1024 as a divisor for each magnitude. In base 10, sizes are shown as powers of 10, using 1000 as a divisor for each magnitude. 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.
- Precision Preference. The number of decimal places (1–10) for display of storage-space sizes.
- Unit Preference. The unit for display of storage-space sizes: Auto, TB, GB, MB. The Auto option lets the system 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, precision is set to 1, and base is set to 10, the size 0.11709 TB is shown as 117.1 GB.
- Temperature Preference. The scale for display of temperature values: Celsius or Fahrenheit.
- Auto Sign Out (minutes). The amount of time that the user's session can be idle before the user is automatically signed out (2–720 minutes).
- Locale. The user's preferred display language, which overrides the system's default display language. Installed language sets include Chinese-Simplified, Chinese-Traditional, Dutch, English, French, German, Italian, Japanese, Korean, and Spanish.

SNMPv3 user accounts have these options:

- User Name.
- Password.
- SNMP User Type. Either: User Access, which allows the user to view the SNMP MIB; or Trap Target, which allows the user to receive SNMP trap notifications. Trap Target uses the IP address set with the Trap Host Address option.
- Authentication Type. Either: MD5 authentication; SHA (Secure Hash Algorithm) authentication; or no authentication. Authentication uses the password set with the Password option.
- Privacy Type. Either: DES (Data Encryption Standard) encryption; AES (Advanced Encryption Standard) encryption; or no encryption. Encryption uses the password set with the Privacy Password option.
- Privacy Password. The encryption password.
- Trap Host Address. The IP address of the host system that will receive SNMP traps.

Table 3 Settings for default users

Name	Password	Roles	Type	Interfaces enabled	Base	Prec.	Units	Temp.	Auto Sign Out	Locale
monitor	!monitor	Monitor	Standard	WBI, CLI	10	1	Auto	Celsius	30 min.	English
manage	!manage	Monitor, Manage		WBI, CLI, FTP, SMI-S						
ftp	!ftp	Monitor, Manage		FTP						

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

Related topics

- ["Configuring user accounts" \(page 44\)](#)

About vdisks

A *vdisk* is a virtual disk that is composed of one or more disks, and has the combined capacity of those disks. The number of disks that a *vdisk* can contain is determined by its RAID level. All disks in a *vdisk* must be the same type (SAS SSD, enterprise SAS, or midline SAS). A maximum of 64 *vdisks* per system can exist.

A *vdisk* can contain different models of disks, and disks with different capacities. If you mix disks with different capacities, the smallest disk determines the logical capacity of all other disks in the *vdisk*, regardless of RAID level. For example, the capacity of a *vdisk* composed of one 500-GB disk and one 750-GB disk is equivalent to a *vdisk* composed of two 500-GB disks.

Each disk has metadata that identifies whether the disk is a member of a *vdisk*, and identifies other members of that *vdisk*. This enables disks to be moved to different slots in a system; an entire *vdisk* to be moved to a different system; and a *vdisk* to be quarantined if disks are detected missing.

In a single-controller system, all *vdisks* are owned by that controller. In a dual-controller system, when a *vdisk* is created the system automatically assigns the owner to balance the number of *vdisks* each controller owns; or, you can select the owner. Typically it does not matter which controller owns a *vdisk*.

In a dual-controller system, when a controller fails, the partner controller assumes temporary ownership of the failed controller's *vdisks* and resources. If a fault-tolerant cabling configuration is used to connect the controllers to drive enclosures and hosts, both controllers' LUNs are accessible through the partner.

When you create a *vdisk* you can use the default chunk size or one that better suits your application. The chunk size is the amount of contiguous data that is written to a disk before moving to the next disk. After a *vdisk* is created its chunk size cannot be changed. For example, if the host is writing data in 16-KB transfers, that size would be a good choice for random transfers because one host read would generate the read of exactly one disk in the volume. That means if the requests are random-like, then the requests would be spread evenly over all of the disks, which is good for performance. If you have 16-KB accesses from the host and a 64-KB block size, then some of the hosts accesses

would hit the same disk; each chunk contains four possible 16-KB groups of data that the host might want to read, which is not an optimal solution. Alternatively, if the host accesses were 128 KB, then each host read would have to access two disks in the vdisk. For random patterns, that ties up twice as many disks.

When you create a vdisk you can also create volumes within it. A volume is a logical subdivision of a vdisk, and can be mapped to controller host ports for access by hosts. The storage system presents only volumes, not vdisks, to hosts.

You can create vdisks with or without volumes by using the Provisioning Wizard, or you can create vdisks manually.

 **TIP:** Best practices for creating vdisks include:

- To maximize capacity, use disks of similar size.
- For greatest reliability, use disks of the same size and rotational speed.
- For storage configurations using many disks, create a few vdisks each containing many disks instead of many vdisks each containing a few disks.
- To maximize capacity and disk usage (but not performance), you can create vdisks larger than 2 TB and divide them into multiple volumes each having a capacity of 2 TB or less. This increases the usable capacity of storage configurations by reducing the total number of parity disks required when using parity-protected RAID levels. This differs from using a *volume* larger than 2 TB, which requires specific support by the host operating system, I/O adapter, and application.
- For maximum use of a dual-controller system's resources, each controller should own a similar number of vdisks.
- Set the chunk size to match the transfer block size of the host application.

Related topics

- ["About RAID levels" \(page 28\)](#)
- ["About spares" \(page 19\)](#)
- ["About volumes" \(page 20\)](#)
- Vdisk topics in ["Provisioning the system" \(page 59\)](#)
- ["Configuring a vdisk" \(page 56\)](#)
- ["Verifying a vdisk" \(page 87\)](#)
- ["Scrubbing a vdisk" \(page 87\)](#)
- Viewing information about a vdisk ([page 100](#)), all vdisks ([page 99](#)), or the system ([page 93](#))
- ["Removing a vdisk from quarantine" \(page 88\)](#)

About spares

A controller automatically reconstructs a fault-tolerant vdisk (RAID 1, 3, 5, 6, 10, 50) when one or more of its disks fails and a compatible spare disk is available. A compatible disk has enough capacity to replace the failed disk and is the same type (SAS SSD, enterprise SAS, or midline SAS).

There are three types of spares:

- *Dedicated spare.* Reserved for use by a specific vdisk to replace a failed disk. Most secure way to provide spares for vdisks but expensive to reserve a spare for each vdisk.
- *Global spare.* Reserved for use by any fault-tolerant vdisk to replace a failed disk.
- *Dynamic spare.* An available compatible disk that is automatically assigned to replace a failed disk in a fault-tolerant vdisk.

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

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

Sparing rules for heterogeneous vdisks

If you upgraded from an earlier release that did not distinguish between enterprise and midline SAS disks, you might have vdisks that contain both types of disks. These are called heterogeneous or mixed vdisks. In the Configuration View panel, the vdisk's RAID-level label includes the suffix `-MIXED`.

For heterogeneous vdisks, the system uses the following logic for global sparing and dynamic sparing. If a vdisk has more than one type of disk in it, the system will look for disks of all types contained in the vdisk. In an effort to migrate heterogeneous vdisks to homogeneous vdisks, the disk type that is most prominent (has the highest number of disks) will be preferred. If all the disk types in a vdisk have the same number of disks, the type that has the smallest capacity disk will be used. If both types have the same capacity disks, enterprise SAS will be the preferred type. Dedicated spares are considered part of a vdisk, so they do not use this logic to choose a preferred disk type since using either type will not change the makeup of the vdisk.

The precedence of spares is as follows:

- Dedicated spares of any type.
- Global spares of preferred type.
- Global spares of non-preferred type.
- Dynamic spares of preferred type (if dynamic sparing is enabled).
- Dynamic spares of non-preferred type (if dynamic sparing is enabled).

Related topics

- ["Configuring dynamic spares" \(page 50\)](#)
- ["Managing dedicated spares" \(page 56\)](#)
- ["Managing global spares" \(page 62\)](#)
- ["Using the Provisioning Wizard" \(page 59\)](#)
- ["Creating a vdisk" \(page 61\)](#)
- Viewing information about a vdisk ([page 100](#)) or all vdisks ([page 99](#))

About volumes

A *volume* is a logical subdivision of a vdisk, and can be mapped to controller host ports for access by hosts. A mapped volume provides the storage for a file system partition you create with your operating system or third-party tools. The storage system presents only volumes, not vdisks, to hosts. A maximum of 128 mappable volumes per vdisk can exist.

You can create a vdisk that has one volume or multiple volumes.

- Single-volume vdisks work well in environments that need one large, fault-tolerant storage space for data on one host. A large database accessed by users on a single host that is used only for that application is an example.
- Multiple-volume vdisks work well when you have very large disks and you want to make the most efficient use of disk space for fault tolerance (parity and spares). For example, you could create one 10-TB RAID-5 vdisk and dedicate one spare to the vdisk. This minimizes the amount of disk space allocated to parity and spares compared to the space required if you created five 2-TB RAID-5 vdisks. However, I/O to multiple volumes in the same vdisk can slow system performance.

When you create volumes you can specify their sizes. If the total size of a vdisk's volumes equals the size of the vdisk, you will not have any free space. Without free space, you cannot add or expand volumes. If you need to add or expand a volume in a vdisk without free space, you can delete a volume to create free space. Or, you can expand the vdisk and then either add a volume or expand a volume to use the new free space.

Volume sizes are aligned to 4-MB boundaries. When a volume is created or expanded, if the resulting size would be less than 4 MB it will be increased to 4 MB; if the resulting size would be greater than 4 MB it will be decreased to the nearest 4-MB boundary. The minimum volume size is 4 MB.

You can use a volume's default name or change it to identify the volume's purpose. For example, a volume used to store payroll information can be named Payroll.

You can create vdisks with volumes by using the Provisioning Wizard, or you can create volumes manually.

Related topics

- ["About vdisks" \(page 18\)](#)
- ["About volume mapping" \(page 22\)](#)
- ["About volume cache options" \(page 22\)](#)
- Volume topics in ["Provisioning the system" \(page 59\)](#)
- ["Changing a volume's name" \(page 58\)](#)
- ["Changing a volume's cache settings" \(page 58\)](#)
- ["Viewing information about a volume" \(page 104\)](#)

About hosts

A *host* identifies an external port that the storage system is attached to. The external port may be a port in an I/O adapter (such as an FC HBA) in a server.

The controllers automatically discover hosts that have sent an `inquiry` command or a `report luns` command to the storage system. Hosts typically do this when they boot up or rescan for devices. When the command from the host occurs, the system saves the host ID. The ID for an FC or SAS host is its WWPN. The ID for an iSCSI host is typically, but not limited to, its IQN. You can also manually create entries for hosts.

You can assign a name to a host to make it easy to recognize for volume mapping. A maximum of 64 names can be assigned.

The Configuration View panel lists hosts by name, or if they are unnamed, by ID.

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 login 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 typically, but not limited to, its IQN. A secret must have 12–16 characters.
- Define CHAP entries in the storage system. If the node name is a host name, then it may be useful to display the hosts that are known to the system.
- Enable CHAP on the storage system. Note that this applies to all iSCSI hosts, in order to avoid security exposures.
- Define the CHAP secret(s) in the host iSCSI initiator.
- Request host login to the storage system. 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 login 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.

Related topics

- ["Using the Configuration Wizard" \(page 37\)](#)
- ["Changing host interface settings" \(page 47\)](#)
- ["Adding a host" \(page 74\)](#)
- ["Removing hosts" \(page 75\)](#)
- ["Changing a host's name" \(page 75\)](#)
- ["Changing host mappings" \(page 75\)](#)
- Viewing information about a host ([page 110](#)) or all hosts ([page 109](#))

About volume mapping


Each volume has default host-access settings that are set when the volume is created; these settings are called the *default mapping*. The default mapping applies to any host that has not been explicitly mapped using different settings. *Explicit mappings* for a volume override its default mapping.

Default mapping enables all attached hosts to see a volume using a specified LUN and access permissions set by the administrator. This means that when the volume is first created, all connected hosts can immediately access the volume using the advertised default mapping settings. This behavior 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.

You can change a volume's default mapping, 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*. You can apply access privileges to one or more of the host ports on either controller. To maximize performance, map a volume to at least one host port on the controller that owns it. To sustain I/O in the event of controller failure, map to at least one host port on each controller.

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 per volume 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.

Volume mapping settings are stored in disk metadata. If enough of the disks used by a volume are moved into a different enclosure, the volume's vdisk can be reconstructed and the mapping data is preserved.

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 vdisk 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.


Related topics

- ["Using the Provisioning Wizard" \(page 59\)](#)
- [Changing a volume's default mapping \(page 66\)](#) or [explicit mappings \(page 67\)](#)
- ["Changing host mappings" \(page 75\)](#)
- [Viewing information about a volume \(page 104\)](#), [snapshot \(page 106\)](#), [host \(page 110\)](#), or [all hosts \(page 109\)](#)

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 a volume's write-back cache setting. *Write-back* is a cache-writing strategy in which the controller receives the data to be written to disks, stores it in the memory buffer, and immediately sends the host operating system a signal that the write operation is complete, without waiting until the data is actually written to the disk. Write-back cache mirrors all of the data from one controller module cache to the other. Write-back cache improves the performance of write operations and the throughput of the controller.

When write-back cache is disabled, *write-through* becomes the cache-writing strategy. Using write-through cache, the controller writes the data to the disks before signaling the host operating system that the process is complete. Write-through cache has lower write 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 super-capacitor technology, if the system loses power, data is not lost. For most applications, this is the preferred setting.

If you are doing random access to this volume, leave the write-back cache enabled.

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

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. Read ahead can be forward (that is, increasing LBAs) or reverse (that is, decreasing LBAs). Increasing the read-ahead cache size can greatly improve performance for multiple sequential read streams. However, increasing read-ahead size will likely decrease random read performance.

- Adaptive. This option works well for most applications: it enables adaptive read-ahead, which allows the controller to dynamically calculate the optimum read-ahead size for the current workload. This is the default.
- Stripe. This option sets the read-ahead size to one stripe. The controllers treat non-RAID and RAID-1 vdisks internally as if they have a stripe size of 512 KB, even though they are not striped.
- Specific size options. These options let you select an amount of data for all accesses.
- Disabled. This 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.

You can also change the optimization mode.

- Standard. This mode works well for typical applications where accesses are a combination of sequential and random; this method is the default. For example, use this mode for transaction-based and database update applications that write small files in random order.
- No-mirror. In this mode each controller stops mirroring its cache metadata to the partner controller. This improves write I/O response time but at the risk of losing data during a failover. ULP behavior is not affected, with the exception that during failover any write data in cache will be lost.
- Atomic write. Not supported in this release.

Related topics

- ["Changing a volume's cache settings" \(page 58\)](#)
- ["Changing system cache settings" \(page 52\)](#)
- ["Viewing information about a volume" \(page 104\)](#)

About managing remote systems

You can add a management object to obtain information from a remote storage system. This allows a local system to track remote systems by their network-port IP addresses and cache their login credentials — the user name and password for a manage-level user on that system. The IP address can then be used in commands that need to interact with the remote system.

After a remote system has been added, you can check the connectivity of host ports in the local system to host ports in that remote system. A port in the local system can only link to ports with the same host interface, such as Fibre Channel (FC), in a remote system.

Communication between local and remote systems is an essential part of the remote replication feature.

Related topics

- ["Adding a remote system" \(page 55\)](#)
- ["Deleting remote systems" \(page 55\)](#)
- ["Viewing information about a remote system" \(page 119\)](#)
- ["Checking links to a remote system" \(page 90\)](#)
- ["About the AssuredRemote replication feature" \(page 121\)](#)

About the Snapshot feature

Snapshot (AssuredSnap™) is a licensed feature that provides data protection by enabling you to create and save snapshots of a volume. Each snapshot preserves the source volume's data state at the point in time when the snapshot was created. Snapshots can be created manually or by using the task scheduler.

When the first snapshot is taken of a standard volume, the system automatically converts the volume into a *master volume* and reserves additional space for snapshot data. This reserved space, called a *snap pool*, stores pointers to the source volume's data. Each master volume has its own snap pool. 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 snapshot's purpose. Any additional unique data written to a snapshot is also stored in the snap pool.

The following figure shows how the data state of a master volume is preserved in the snap pool by two snapshots taken at different points in time. The dotted line used for the snapshot borders indicates that snapshots are logical volumes, not physical volumes as are master volumes and snap pools.




Figure 1 Relationship between a master volume and its snapshots and snap pool

The snapshot feature uses the single copy-on-write method to capture only data that has changed. That is, if a block is to be overwritten on the master volume, and a snapshot depends on the existing data in the block being overwritten, the data is copied from the master volume to the snap pool before the data is changed. All snapshots that depend on the older data are able to access it from the same location in the snap pool; this reduces the impact of snapshots when writing to a master volume. In addition, only a single copy-on-write operation is performed on the master volume.

The storage system allows a maximum number of snapshots to be retained, as determined by an installed license. For example, if your license allows four snapshots, when the fifth snapshot is taken an error message informs you that you have reached the maximum number of snapshots allowed on your system. Before you can create a new snapshot you must either delete an existing snapshot, or purchase and install a license that increases the maximum number of snapshots.

The snapshot service has two features for reverting data back to original data:

- Deleting only modified data on a snapshot. For snapshots that have been made accessible as read-write, you can delete just the modified (write) data that was written directly to a snapshot. When the modified data is deleted, the snapshot data reverts to the original data that was snapped. This feature is useful for testing an application, for example. You might want to test some code, which writes data to the snapshot. Rather than having to take another snapshot, you can just delete any write data and start again.
- Rolling back the data in a source volume. The rollback feature enables you to revert the data in a source volume to the data that existed when a specified snapshot was created (preserved data). Alternatively, the rollback can include data that has been modified (write data) on the snapshot since the snapshot was taken. For example, you might want to take a snapshot, mount/present/map that snapshot for read/write, and then install new software on that snapshot for test purposes. If the software installation is successful, you can rollback the master volume to the contents of the modified snapshot (preserved data plus the write data).

The following figure shows the difference between rolling back the master volume to the data that existed when a specified snapshot was created (preserved), and rolling back preserved and modified data.

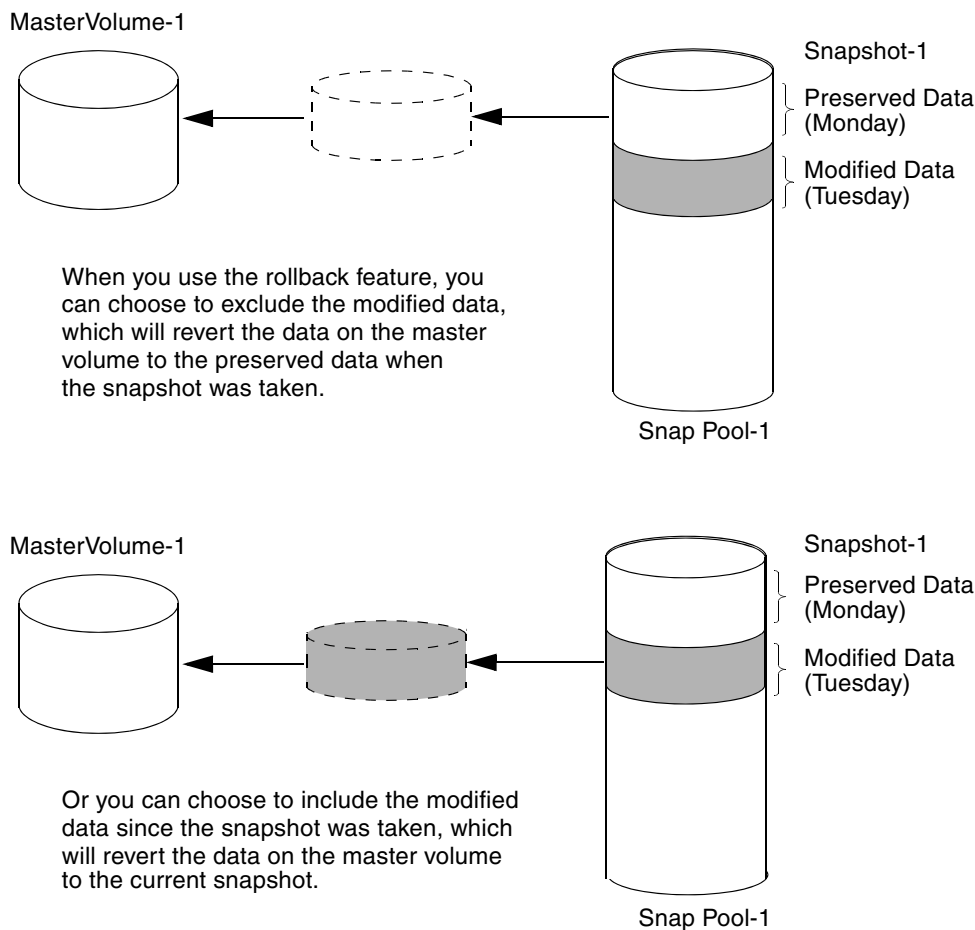


Figure 2 Rolling back a master volume

Snapshot operations are I/O-intensive. Every write to a unique location in a master volume after a snapshot is taken will cause an internal read and write operation to occur in order to preserve the snapshot data. If you intend to create snapshots of, create volume copies of, or replicate volumes in a vdisk, ensure that the vdisk contains no more than four master volumes, snap pools, or both. For example: 2 master volumes and 2 snap pools; 3 master volumes and 1 snap pool; 4 master volumes and 0 snap pools.

Related topics

- ["Installing a license" \(page 41\)](#)
- [Creating a snapshot \(page 69\)](#) or [multiple snapshots \(page 68\)](#)
- [Changing a snapshot's default mapping \(page 66\)](#) or [explicit mappings \(page 67\)](#)
- ["Deleting snapshots" \(page 70\)](#)
- ["Resetting a snapshot" \(page 70\)](#)
- [Viewing information about a snapshot \(page 106\)](#), [a vdisk \(page 100\)](#), [all vdisks \(page 99\)](#), or [the system \(page 93\)](#)
- ["Rolling back a volume" \(page 73\)](#)
- ["Deleting schedules" \(page 78\)](#)

About the Volume Copy feature

Volume Copy (AssuredCopy™) is a licensed feature that enables you to copy a volume or a snapshot to a new standard volume.


While a snapshot is a point-in-time logical copy of a volume, the volume copy service creates a complete "physical" copy of a volume within a storage system. It is an exact copy of a source volume as it existed at the time the volume copy operation was initiated, consumes the same amount of space as the source volume, and is independent from an I/O perspective. Volume independence is a key distinction of a volume copy (versus a snapshot, which is a "virtual" copy and dependent on the source volume).

Benefits include:

- **Additional data protection.** An independent copy of a volume (versus logical copy through snapshot) provides additional data protection against a complete master volume failure. If the source master volume fails, the volume copy can be used to restore the volume to the point in time the volume copy was taken.
- **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 snapshot) so that I/O is to each set of blocks respectively; application I/O transactions are not competing with each other when accessing the same data blocks.


The following figure illustrates how volume copies are created.

Creating a volume copy from a standard or master volume



1. Volume copy request is made with a standard volume or a master volume as the source.
2. If the source is a standard volume, it is converted to a master volume and a snap pool is created.
3. A new volume is created for the volume copy, and a hidden, transient snapshot is created.
4. Data is transferred from the transient snapshot to the new volume.
5. On completion, the transient volume is deleted and the new volume is a completely independent copy of the master volume, representing the data that was present when the volume copy was started.

Creating a volume copy from a snapshot



1. A master volume exists with one or more snapshots associated with it. Snapshots can be in their original state or they can be modified.
2. You can select any snapshot to copy, and you can specify that the modified or unmodified data be copied.
3. On completion, the new volume is a completely independent copy of the snapshot. The snapshot remains, though you can choose to delete it.

Figure 3 Creating a volume copy from a master volume or a snapshot

Snapshot operations are I/O-intensive. Every write to a unique location in a master volume after a snapshot is taken will cause an internal read and write operation to occur in order to preserve the snapshot data. If you intend to create snapshots of, create volume copies of, or replicate volumes in a vdisk, ensure that the vdisk contains no more than four master volumes, snap pools, or both. For example: 2 master volumes and 2 snap pools; 3 master volumes and 1 snap pool; 4 master volumes and 0 snap pools.

Guidelines to keep in mind when performing a volume copy include:

- The destination vdisk must be owned by the same controller as the source volume.
- The destination vdisk must have free space that is at least as large as the amount of space allocated to the original volume. A new volume will be created using this free space for the volume copy.
- The destination vdisk does not need to have the same attributes (such as disk type, RAID level) as the volume being copied.
- Once the copy is complete, the new volume will no longer have any ties to the original.
- Volume Copy makes a copy from a snapshot of the source volume; therefore, the snap pool for the source volume must have sufficient space to store snapshot data when performing this copy.

Related topics

- ["Installing a license" \(page 41\)](#)
- ["Creating a volume copy" \(page 71\)](#)
- ["Aborting a volume copy" \(page 73\)](#)
- ["Viewing information about a volume" \(page 104\)](#)
- ["Deleting schedules" \(page 78\)](#)

About the AssuredRemote replication feature

See ["Using AssuredRemote to replicate volumes" \(page 121\)](#).

About the VDS and VSS hardware providers

Virtual Disk Service (VDS) enables host-based applications to manage vdisks and volumes. Volume Shadow Copy Service (VSS) enables host-based applications to manage snapshots. A license is required to enable VDS and VSS hardware providers, so hosts can manage vdisks, volumes, and snapshots in the storage system. For more information, see the VDS and VSS hardware provider documentation for your product.


About the Storage Replication Adapter (SRA)

The SRA is a host-software component installed on a Microsoft Windows Server operating system that allows VMware vCenter Site Recovery Manager software to control certain aspects of the replication feature in storage systems. The presence of the SRA allows the disaster recovery software to automatically coordinate virtual-machine failover and failback between a protected data center and a disaster recovery site. A license is required to enable the SRA.

About RAID levels

The RAID controllers enable you to set up and manage vdisks, whose storage may be spread across multiple disks. This is accomplished through firmware resident in the RAID controller. RAID refers to vdisks 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 vdisk fails.

Hosts see each partition of a vdisk, known as a volume, as a single disk. A volume is actually a portion of the storage space on disks behind a RAID controller. The RAID controller firmware makes each volume appear as one very large disk. Depending on the RAID level used for a vdisk, the disk presented to hosts has advantages in fault-tolerance, cost, performance, or a combination of these.

 **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
 - Describe the expansion capability for different RAID levels
-

NOTE: To create an NRAID, RAID-0, or RAID-3 vdisk, you must use the CLI.

Table 4 Example applications and RAID levels

Application	RAID level
Testing multiple operating systems or software development (where redundancy is not an issue)	NRAID
Fast temporary storage or scratch disks for graphics, page layout, and image rendering	0
Workgroup servers	1 or 10
Video editing and production	3
Network operating system, databases, high availability applications, workgroup servers	5
Very large databases, web server, video on demand	50
Mission-critical environments that demand high availability and use large sequential workloads	6

Table 5 RAID level comparison

RAID level	Min. disks	Description	Strengths	Weaknesses
NRAID	1	Non-RAID, nonstriped mapping to a single disk	Ability to use a single disk to store additional data	Not protected, lower performance (not striped)
0	2	Data striping without redundancy	Highest performance	No data protection: if one disk fails all data is lost
1	2	Disk mirroring	Very high performance and data protection; minimal penalty on write performance; protects against single disk failure	High redundancy cost overhead: because all data is duplicated, twice the storage capacity is required
3	3	Block-level data striping with dedicated parity disk	Excellent performance for large, sequential data requests (fast read); protects against single disk failure	Not well-suited for transaction-oriented network applications; write performance is lower on short writes (less than 1 stripe)
5	3	Block-level data striping with distributed parity	Best cost/performance for transaction-oriented networks; very high performance and data protection; supports multiple simultaneous reads and writes; can also be optimized for large, sequential requests; protects against single disk failure	Write performance is slower than RAID 0 or RAID 1
6	4	Block-level data striping with double distributed parity	Best suited for large sequential workloads; non-sequential read and sequential read/write performance is comparable to RAID 5; protects against dual disk failure	Higher redundancy cost than RAID 5 because the parity overhead is twice that of RAID 5; not well-suited for transaction-oriented network applications; non-sequential write performance is slower than RAID 5
10 (1+0)	4	Stripes data across multiple RAID-1 sub-vdisks	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
50 (5+0)	6	Stripes data across multiple RAID-5 sub-vdisks	Better random read and write performance and data protection than RAID 5; supports more disks than RAID 5; protects against multiple disk failures	Lower storage capacity than RAID 5

Table 6 Vdisk expansion by RAID level

RAID level	Expansion capability	Maximum disks
NRAID	Cannot expand.	1
0, 3, 5, 6	You can add 1–4 disks at a time.	16
1	Cannot expand.	2
10	You can add 2 or 4 disks at a time.	16
50	You can add one sub-vdisk at a time. The added sub-vdisk must contain the same number of disks as each of the existing sub-vdisks.	32

About size representations

Parameters such as names of users and volumes have a maximum length in bytes. ASCII characters are 1 byte; most Latin (Western European) characters with diacritics are 2 bytes; most Asian characters are 3 bytes.

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 RAIDar, the base for entry and display of storage-space sizes can be set per user or per session. When entering storage-space sizes only, either base-2 or base-10 units can be specified.

Table 7 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 8 Decimal (radix) point character by locale

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

Related topics

- ["About user accounts" \(page 17\)](#)

About the system date and time

You can change the storage system's date and time, which are displayed in the System Status panel. It is important to set the date and time so that entries in system logs and event-notification email messages have correct time stamps.

You can set the date and time manually or configure the system to use Network Time Protocol (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 Coordinated Universal Time (UTC), which provides several options:

- If you want to synchronize the times and logs between storage devices installed in multiple time zones, set all the storage devices to use UTC.
- If you want 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 U.S. daylight savings time. You must make such adjustments manually.







Related topics

- ["Changing the system date and time" \(page 47\)](#)

About storage-space color codes

RAIDar panels use the following color codes to identify how storage space is used.



















Table 9 Storage-space color codes

Area	Color	Meaning
Overview panels		Total space
		Available/free space
		Used space
		Reserved/overhead space, used for parity and snap pools, for example
Vdisk panels		Space used by spares
		Wasted space, due to use of mixed disk sizes

About Configuration View icons

The Configuration View panel uses the following icons to let you view physical and logical components of the storage system.

Table 10 Configuration View icons

Icon	Meaning	Icon	Meaning
	Show all subcomponents		Snapshot
	Hide all subcomponents		Snap pool
	Show the component's subcomponents		Replication-prepared volume
	Hide the component's subcomponents		Local primary volume
	Storage system		Local secondary volume
	Enclosure		Local replication image
	Host/initiator		Remote primary volume
	Vdisk		Remote secondary volume
	Standard or master volume		Remote replication image

About disk failure and vdisk reconstruction

Vdisk reconstruction does not require I/O to be stopped, so the vdisk can continue to be used while the Reconstruct utility runs. Vdisk reconstruction starts automatically when all of the following are true:

- One or more disks fail in a fault-tolerant vdisk (RAID 1, 3, 5, 6, 10, or 50)
- The vdisk is still operational
- Compatible spares are available

The storage system automatically uses the spares to reconstruct the vdisk. A compatible spare has a capacity equal to or greater than the smallest disk in the vdisk, has enough capacity to replace a failed disk, and is the same type (SAS SSD, enterprise SAS, or midline SAS) as those disks. If no compatible spares are available, reconstruction does not start automatically. To start reconstruction manually, replace each failed disk and then do one of the following:

- Add each new disk as either a dedicated spare or a global spare. Remember that a global spare might be taken by a different critical vdisk than the one you intended. When a global spare replaces a disk in a vdisk, the global spare's icon in the enclosure view changes to match the other disks in that vdisk.
- Enable the Dynamic Spare Capability option to use the new disks without designating them as spares.
- Change a dedicated spare from a different vdisk to either a global spare or a dedicated spare for the degraded vdisk.

RAID-6 reconstruction behaves as follows:

- During online initialization, if one disk fails, initialization continues and the resulting vdisk will be degraded (FTDN status). After initialization completes, the system can use a compatible spare to reconstruct the vdisk.
- During online initialization, if two disks fail, initialization stops (CRIT status). The system can use two compatible spares to reconstruct the vdisk.
- During vdisk operation, if one disk fails and a compatible spare is available, the system begins to use that spare to reconstruct the vdisk. If a second disk fails during reconstruction, reconstruction continues until it is complete, regardless of whether a second spare is available. If the spare fails during reconstruction, reconstruction stops.
- During vdisk operation, if two disks fail and only one compatible spare is available, the system waits five minutes for a second spare to become available. After five minutes, the system begins to use that spare to reconstruct one disk in the vdisk (referred to as “fail 2, fix 1” mode). If the spare fails during reconstruction, reconstruction stops.
- During vdisk operation, if two disks fail and two compatible spares are available, the system uses both spares to reconstruct the vdisk. If one of the spares fails during reconstruction, reconstruction proceeds in “fail 2, fix 1” mode. If the second spare fails during reconstruction, reconstruction stops.

When a disk fails, its Fault LED illuminates amber. When a spare is used as a reconstruction target, its Activity LED blinks green. For details about LED states, see your product’s Setup Guide.

NOTE: Reconstruction can take hours or days to complete, depending on the vdisk RAID level and size, disk speed, utility priority, and other processes running on the storage system. You can stop reconstruction only by deleting the vdisk.

About data protection in a single-controller storage system

An AssuredSAN 4004 Series storage system can be purchased or operated with a single controller. Because single-controller mode is not a redundant configuration, this section presents some considerations concerning data protection.

A volume’s default caching mode 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 target volume’s enclosure is powered off without a proper shut down. Data remains in the controller’s cache and associated volumes will be missing that data. This can result in data loss or in some cases volume loss; for example, if using snapshot functionality a snap pool might become inaccessible and the master volume could go offline.

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

To avoid the possibility of data loss in case the controller fails you can change a volume’s caching mode 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 details about caching modes see ["About volume cache options" \(page 22\)](#). To change a volume’s caching mode, see ["Changing a volume’s cache settings" \(page 58\)](#).

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, SNMP, or SMI-S 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 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 RAIDar's Save Logs panel or the FTP interface's `get_logs` command. 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 panel, see ["Saving logs" \(page 83\)](#). For information about using the FTP interface, see the RAIDar User Guide.

Related topics

- ["Configuring email notification" \(page 43\)](#) (for push mode)
- ["Configuring SNMP notification" \(page 44\)](#) (for pull mode)
- ["Changing management interface settings" \(page 42\)](#) (to enable SNMP or SMI-S for pull mode)
- ["Enabling/disabling managed logs" \(page 55\)](#)
- ["Testing notifications" \(page 86\)](#)

About performance monitoring

The storage system samples disk-performance statistics every quarter hour and retains performance data for 6 months. You can view these historical performance statistics to identify disks that are experiencing errors or are performing poorly.

RAIDar displays historical performance statistics in graphs for ease of analysis. You can view historical performance statistics either for a single disk or for all disks in a vdisk. By default, the graphs will show the latest 50 data samples, but you can specify the time period to display. If the specified time period includes more than 50 samples, their data will be aggregated into 50 samples; the graphs show a maximum of 50 samples. Data shown will be up-to-date as of the time it is requested for display, and summary statistics will be updated when a new sample is taken.

Disk-performance graphs include:

- Data Transferred
- Data Throughput
- I/O
- IOPS
- Average Response Time
- Average I/O Size
- Disk Error Counters
- Average Queue Depth

Vdisk-performance graphs include:

- Data Transferred
- Data Throughput
- Average Response Time

You can save historical statistics in CSV format to a file for import into a spreadsheet or other third-party application. You can also reset historical statistics, which clears the retained data and continues to gather new samples.

NOTE: RAIDar does not show live statistics. For information about viewing live statistics, see the CLI Reference Guide.

Related topics

- ["Vdisk performance" \(page 101\)](#) (to view historical performance statistics for a vdisk)
- ["Disk performance" \(page 113\)](#) (to view historical performance statistics for a disk)
- ["Resetting historical disk-performance statistics" \(page 90\)](#)
- ["Saving historical disk-performance statistics" \(page 91\)](#)

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. The controller-module firmware-update algorithm supports the following scenarios for a dual-controller system:

- 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.

When you update controller firmware, you can use the Partner Firmware Update option to ensure that the same firmware version is installed in both controller modules. This option 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 in that controller is transferred to the partner controller.
- 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 newer firmware version in either controller is transferred to the other controller.

For information about the procedures to update firmware in controller modules, expansion modules, and disk drives, see ["Updating firmware" \(page 79\)](#). That topic also describes how to use the activity progress interface to view detailed information about the progress of a firmware-update operation.

2 Configuring the system

Using the Configuration Wizard

The Configuration Wizard helps you initially configure the system or change system configuration settings.

The wizard guides you through the following steps. For each step you can view help by clicking the help icon  in the wizard panel. As you complete steps they are highlighted at the bottom of the panel. If you cancel the wizard at any point, no changes are made.

- Change passwords for the default users
- 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
- Confirm changes and apply them

When you complete this wizard you are given the option to start the Provisioning Wizard to provision storage.

Step 1: Starting the wizard

1. In the Configuration View panel, right-click the system and select either **Configuration > Configuration Wizard** or **Wizards > Configuration Wizard**. The wizard panel appears.
2. Click **Next** to continue.

Step 2: Changing default passwords

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

1. To secure the storage system, set a new password for each default user. A password is case sensitive and can include a maximum of 32 bytes using characters except a backslash, comma, angle bracket, or double quote.
2. Click **Next** to continue.

Step 3: Configuring network ports

You can configure addressing parameters for each controller's network port. You can set static IP values or use DHCP.

In DHCP mode, network port IP address, subnet mask, and gateway values are obtained 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:

- DHCP: disabled
- 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, the following initial values are set and remain set until the system is able to contact a DHCP server for new addresses:

- Controller IP addresses: 169.254.*x*.*x* (where the value of *x*.*x* is the lowest 16 bits of the controller serial number)
- IP subnet mask: 255.255.0.0
- Gateway IP address: 0.0.0.0

169.254.*x*.*x* addresses (including gateway 169.254.0.1) are on a private subnet that is reserved for unconfigured systems and the addresses are not routable. This prevents the DHCP server from reassigning the addresses and

possibly causing a conflict where two controllers have the same IP address. As soon as possible, change these IP values to proper values for your network.

△ **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. Set the IP address source to **DHCP**.
2. Click **Next** to continue.

To set static IP values for network ports

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

Step 4: Enabling system-management services

You can enable or disable management interfaces to limit the ways in which users and host-based management applications can access the storage system. Network management interfaces operate out-of-band and do not affect host I/O to the system. The network options are:

- Web Browser Interface (WBI). The primary interface for managing the system. You can enable use of HTTP, or HTTPS for increased security, or both.
- Command Line Interface (CLI). An advanced user interface for managing the system. You can enable use of Telnet, or SSH (secure shell) for increased security, or both.
- Storage Management Initiative Specification (SMI-S). Used for management of the system through your network. You can enable use of unencrypted or encrypted SMI-S:
 - Enable. Selecting this option allows SMI-S clients to communicate with each controller's embedded SMI-S provider via HTTP port 5988. HTTP port 5988 and HTTPS port 5989 cannot be enabled at the same time, so enabling this option will disable port 5989.
 - Encrypted. Additionally selecting this option allows SMI-S clients to communicate with each controller's embedded SMI-S provider via HTTPS port 5989. HTTPS port 5989 and HTTP port 5988 cannot be enabled at the same time, so enabling this option will disable port 5988. This is the default.

The Storage Management Initiative Specification (SMI-S) is a Storage Networking Industry Association (SNIA) standard that enables interoperable management for storage networks and storage devices. SMI-S replaces multiple disparate managed object models, protocols, and transports with a single object-oriented model for each type of component in a storage network. The specification was created by SNIA to standardize storage management solutions. SMI-S enables management applications to support storage devices from multiple vendors quickly and reliably because they are no longer proprietary. SMI-S detects and manages storage elements by type, not by vendor.

- File Transfer Protocol (FTP). A secondary interface for installing firmware updates, downloading logs, and installing a license.
- Simple Network Management Protocol (SNMP). Used for 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. This is disabled by default.
- 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 management interfaces operate through the data path and can slightly reduce I/O performance. The in-band option is:

- In-band SES Capability. Used for in-band monitoring of system status based on SCSI Enclosure Services (SES) data.

If a service is disabled, it continues to run but cannot be accessed. To allow specific users to access WBI, CLI, FTP or SMI-S, see ["About user accounts" \(page 17\)](#).

To change management interface settings

1. Enable the options that you want to use to manage the storage system, and disable the others.
2. Click **Next** to continue.

Step 5: Setting system information

Enter a name, contact person, location, and description for the system. Each value can include a maximum of 79 bytes, using characters except angle bracket, double quote, or backslash. 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.

Click **Next** to continue.

Step 6: Configuring event notification

Configure email addresses and SNMP trap hosts to receive event notifications, and configure the managed logs feature.

1. In the Event Notifications section, set the options:
 - 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). The default is none, which disables email notification.
 - SMTP Server address. The IP address of the SMTP mail server to use for the email messages. If the mail server is not on the local network, make sure that the gateway IP address was set in the network configuration step.
 - Sender Name. The sender name that is 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. For example: `Storage-1`. If no sender name is set, a default name is created.
 - Sender Domain. The domain name that is 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. For example: `MyDomain.com`. If the domain name is not valid, some email servers will not process the mail. The default is `mydomain.com`.
 - Email Address fields. Up to three email addresses that the system should send notifications to. Email addresses 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`.
2. In the SNMP Configuration section, set the options:
 - 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). The default is none, which disables SNMP notification.
 - Read Community. The SNMP read password for your network. This password is also included in traps that are sent. The value is case sensitive; can include any character except angle brackets, single quote, and double quote; and can have a maximum of 31 bytes. The default is `public`.
 - Write Community. The SNMP write password for your network. The value is case sensitive; can include any character except angle brackets, single quote, and double quote; and can have a maximum of 31 bytes. The default is `private`.
 - Trap Host Address fields. IP addresses of up to three host systems that are configured to receive SNMP traps.
3. In the Managed Logs Notifications section, set the options:
 - Log Destination. The email address of the log-collection system. The email addresses must use the format `user-name@domain-name` and can have a maximum of 320 bytes. For example: `LogCollector@MyDomain.com`.
 - **Include Logs**. When the managed logs feature is enabled, this option activates "push" mode, which automatically attaches system log files to managed-logs email notifications that are sent to the log-collection system.

NOTE: These options configure the managed logs feature but do not enable it, which is done on the Configuration > Advanced Settings > System Utilities panel.

4. Click **Next** to continue.

Step 7: Configuring host ports

To enable the system to communicate with hosts or with remote systems having FC or iSCSI interfaces, you must configure host-interface options. For SAS ports there are no host-interface options. Click **Next** to continue.

AssuredSAN 4004 Series CNC controller module host ports can be configured as FC or iSCSI ports. For both controller modules, all ports can be set to FC, all ports can be set to iSCSI, or the first two ports in each controller module can be set to FC and the second two ports can be set to iSCSI. FC ports support use of 16-Gbit/s or 8-Gbit/s SFPs. A 16-Gbit/s SFP can run at 16 Gbit/s, 8 Gb/s, or 4 Gb/s or auto-negotiate its link speed. An 8-Gbit/s SFP can run at 8 Gb/s or 4 Gb/s or auto-negotiate its link speed.

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

To change FC host interface settings

1. Set the Speed option to the proper value to communicate with the host. The speed can be set to auto (the default), which auto-negotiates the proper link speed with the host, or to 4Gb, 8Gb, or 16Gb (Gbit/s). 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. Click **Next** to continue.

To change iSCSI host interface settings

1. In the Common Settings for iSCSI section of the panel, set the options that apply to all iSCSI ports:
 - Authentication (CHAP). Enables or disables use of Challenge Handshake Authentication Protocol. Disabled by default.

NOTE: CHAP records for iSCSI login authentication must be defined if CHAP is enabled. To create CHAP records, see "[Configuring CHAP](#)" (page 76).

- Jumbo Frames. Enables or disables support for jumbo frames. A normal frame can contain 1500 bytes whereas a jumbo frame can contain a maximum of 9000 bytes for larger data transfers. Disabled by default.

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

- iSCSI IP Version. Specifies whether IP values use Internet Protocol version 4 (IPv4) or version 6 (IPv6) format. IPv4 uses 32-bit addresses. IPv6 uses 128-bit addresses. The default is IPv4.
 - iSNS. Enables or disables registration with a specified Internet Storage Name Service server, which provides name-to-IP-address mapping. Disabled by default.
 - iSNS Address. Specifies the IP address of an iSNS server. The default address is all zeroes.
 - Alternate iSNS Address. Specifies the IP address of an alternate iSNS server, which can be on a different subnet. The default address is all zeroes.
2. In the upper section of the panel, set the port-specific options:
 - 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 CNC 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
- Netmask. For IPv4, subnet mask for assigned port IP address. The default is 255.255.255.0.
- Gateway. For IPv4, gateway IP address for assigned port IP address. The default is 0.0.0.0.
- Default Router. For IPv6, default router for assigned port IP address. If the gateway was set for IPv4 and then ports were switched to IPv6, the default is `:: IPv4-address`; otherwise, the default is `::` (the short form of all zeroes).
- Link-Local Address. For IPv6, the link-local address that is automatically generated from the MAC address and assigned to the port.

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

3. Click **Next** to continue.

Step 8: Confirming configuration changes

Confirm that the values 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 setting changes and finish the wizard.

Installing a license

A license is required to expand Snapshot limits and to use Volume Copy, Replication, VDS, VSS, and the Storage Replication Adapter (SRA). The license is specific to a controller enclosure serial number and firmware version.

If a permanent license is not installed and you want to try these features before buying a permanent license, you can create a temporary license one time. A temporary license will expire 60 days from the time it is created. After creating a temporary license, each time you sign in to RAIDar, a message specifies the time remaining in the trial period. If you do not install a permanent license before the temporary license expires, you cannot create new items with these features; however, you can continue to use existing snapshots and create new snapshots up to the base maximum limit.

After a temporary license is created or a permanent license is installed, the option to create a temporary license is no longer displayed.

To view information about system licenses

In the Configuration View panel, right-click the system and select **Tools > Install License**.

The System Licenses table shows the following information about licensed features:

- Feature. The name of the licensed feature.
- Base. Either:
 - The number of components that users can create without a license.
 - N/A. Not applicable.
- License. Either:
 - The number of user-created components that the installed license supports.
 - Enabled or Disabled.
- In Use. Either:
 - The number of user-created components that exist.
 - N/A. Not applicable.
- Max Licensable. Either:
 - The number of user-created components that the maximum license supports.
 - N/A. Not applicable.

- Expiration. One of the following:
 - Never. License is purchased and doesn't expire.
 - Number of days remaining for a temporary license.
 - Expired. Temporary license has expired and cannot be renewed.
 - Expired/Renewable. Temporary license has expired and can be renewed.
 - N/A. No license installed.

The panel also shows the licensing serial number (controller enclosure serial number) and licensing version number (controller firmware version), for which a license file must be generated in order to successfully install.

To create a temporary license

1. In the Configuration View panel, right-click the system and select **Tools > Install License**. If the option to create a temporary license is available, the End User License Agreement appears in the lower portion of the license panel.
2. Read the license agreement.
3. If you accept the terms of the license agreement, select the checkbox. A confirmation dialog appears.
4. Click **Yes** to start the trial period. The feature's Expiration value shows the number of days remaining in the trial period; the trial period will expire on the last day. When the trial period expires, the value changes to Expired or Expired/Renewable.

To install a permanent license

1. Ensure that:
 - The license file is saved to a network location that you can access from RAIDar.
 - You are signed into the controller enclosure that the file was generated for.
2. In the Configuration View panel, right-click the system and select **Tools > Install License**.
3. Click **Browse** to locate and select the license file.
4. Click **Install License File**. If installation succeeds, the System Licenses table is updated. The licensing change takes effect immediately. The feature's Expiration value shows Never.

Configuring system services

Changing management interface settings

You can enable or disable management interfaces to limit the ways in which users and host-based management applications can access the storage system. Network management interfaces operate out-of-band and do not affect host I/O to the system. The network options are:

- Web Browser Interface (WBI). The primary interface for managing the system. You can enable use of HTTP, of HTTPS for increased security, or both.
- Command Line Interface (CLI). An advanced user interface for managing the system. You can enable use of Telnet, of SSH (secure shell) for increased security, or both.
- Storage Management Initiative Specification (SMI-S). Used for management of the system through your network. You can enable use of unencrypted or encrypted SMI-S:
 - Enable. Selecting this option allows SMI-S clients to communicate with each controller's embedded SMI-S provider via HTTP port 5988. HTTP port 5988 and HTTPS port 5989 cannot be enabled at the same time, so enabling this option will disable port 5989.
 - Encrypted. Additionally selecting this option allows SMI-S clients to communicate with each controller's embedded SMI-S provider via HTTPS port 5989. HTTPS port 5989 and HTTP port 5988 cannot be enabled at the same time, so enabling this option will disable port 5988. This is the default.
- 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. This is disabled by default.
- 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 management interfaces operate through the data path and can slightly reduce I/O performance. The in-band option is:

- In-band SES Capability. Used for in-band monitoring of system status based on SCSI Enclosure Services (SES) data.

If a service is disabled, it continues to run but cannot be accessed. To allow specific users to access WBI, CLI, or FTP, see ["About user accounts" \(page 17\)](#).

To change management interface settings

1. In the Configuration View panel, right-click the system and select **Configuration > Services > Management**.
2. Enable the options that you want to use to manage the storage system, and disable the others.
3. Click **Apply**. If you disabled any options, a confirmation dialog appears.
4. Click **Yes** to continue; otherwise, click **No**. If you clicked Yes, a processing dialog appears. When processing is complete a success dialog appears.
5. Click **OK**.

Configuring email notification

You can configure email-notification settings for events and managed logs. For an overview of the managed logs feature, see ["About managed logs" \(page 32\)](#).

To configure email notification for events

1. In the Configuration View panel, right-click the system and select **Configuration > Services > Email Notification**.
2. In the main panel, set the options:
 - 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). The default is none, which disables email notification.
 - SMTP Server address. The IP address of the SMTP mail server to use for the email messages. If the mail server is not on the local network, make sure that the gateway IP address is set in System Settings > Network Interfaces.
 - Sender Name. The sender name that is joined with an @ symbol to the domain name to form the "from" address for 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. For example: `Storage-1`. If no sender name is set, a default name is created.
 - Sender Domain. The domain name that is joined with an @ symbol to the sender name to form the "from" address for 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. For example: `MyDomain.com`. If the domain name is not valid, some email servers will not process the mail. The default is `mydomain.com`.
 - Email Address fields. Up to three email addresses that the system should send notifications to. Email addresses 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`.
3. Click **Apply**.
4. Send a test message to the configured destinations as described on [page 86](#).

To configure email notification for managed logs

1. In the Configuration View panel, right-click the system and select **Configuration > Services > Email Notification**.
2. In the main panel, set the options:
 - Log Destination. The email address of the log-collection system. The email addresses must use the format `user-name@domain-name` and can have a maximum of 320 bytes. For example: `LogCollector@MyDomain.com`.
 - **Include Logs**. When the managed logs feature is enabled, this option activates "push" mode, which automatically attaches system log files to managed-logs email notifications that are sent to the log-collection system. This option is disabled by default.

3. Click **Apply**.
4. Enable log management as described on [page 55](#).
5. Send a test message to the configured destination as described on [page 86](#).

Configuring SNMP notification

To configure SNMP notification of events

1. In the Configuration View panel, right-click the system and select **Configuration > Services > SNMP Notification**.
2. In the main panel, set the options:
 - 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). The default is none, which disables SNMP notification.
 - Read Community. The SNMP read password for your network. This password is also included in traps that are sent. The value is case sensitive; can include any character except angle brackets, single quote, and double quote; and can have a maximum of 31 bytes. The default is `public`.
 - Write Community. The SNMP write password for your network. The value is case sensitive; can include any character except angle brackets, single quote, and double quote; and can have a maximum of 31 bytes. The default is `private`.
 - Trap Host Address fields. IP addresses of up to three host systems that are configured to receive SNMP traps.
3. Click **Apply**.
4. Optionally, send a test message to the configured destinations as described on [page 86](#).

Configuring user accounts

Adding users

You can create either a general user that can access the WBI (RAIDar), CLI, FTP or SMI-S interfaces, or an SNMPv3 user that can access the MIB or receive trap notifications. SNMPv3 user accounts support SNMPv3 security features such as authentication and encryption.

To add a general user

1. In the Configuration View panel, right-click the system and select **Configuration > Users > Add New User**.
2. In the main panel, set the options:
 - User Name. A user name is case sensitive; cannot already exist in the system; cannot include a comma, double quote, backslash, or space; and can have a maximum of 29 bytes.
 - Password. A password is case sensitive and can include a maximum of 32 bytes using characters except a backslash, comma, angle bracket, or double quote.
 - Select **Standard User**.
 - User Roles. Select Monitor to let the user view system settings, or Manage to let the user view and change system settings. You cannot change the roles of user manage.
 - User Type. Select an option to identify the user's experience level: Standard, Advanced, or Diagnostic. This parameter does not affect access to commands.
 - **WBI Access**. Allows access to the web-based management interface.
 - **CLI Access**. Allows access to the command-line management interface.
 - **FTP Access**. Allows access to the FTP interface, which can be used instead of the WBI to install firmware updates and download logs.
 - **SMI-S Access**. Allows access to the SMI-S interface, used for management of the system through your network.
 - Base Preference. Select the base for entry and display of storage-space sizes, either Base 10 or Base 2. In base 2, sizes are shown as powers of 2, using 1024 as a divisor for each magnitude. In base 10, sizes are shown as powers of 10, using 1000 as a divisor for each magnitude. 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.
 - Precision Preference. Select the number of decimal places (1–10) for display of storage-space sizes.

- **Unit Preference.** Select a unit for display of storage-space sizes: Auto, TB, GB, MB. The Auto option lets the system 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, precision is set to 1, and base is set to 10, the size 0.11709 TB is shown as 117.1 GB. The default is Auto.
- **Temperature Preference.** Specifies the scale to use for temperature values: Celsius or Fahrenheit.
- **Auto Sign Out (minutes).** Select the amount of time that the user's session can be idle before the user is automatically signed out (2–720 minutes). The default is 30 minutes.
- **Locale.** The user's preferred display language, which overrides the system's default display language. Installed language sets include Chinese-Simplified, Chinese-Traditional, Dutch, English, French, German, Italian, Japanese, Korean, and Spanish.

3. Click **Add User**.

To add an SNMPv3 user

1. In the Configuration View panel, right-click the system and select **Configuration > Users > Add New User**.
2. In the main panel, set the options:
 - **User Name.** A user name is case sensitive; cannot already exist in the system; cannot include a comma, double quote, backslash, or space; and can have a maximum of 29 bytes.
 - **Password.** A password is case sensitive and can include a maximum of 32 bytes using characters except a backslash, comma, angle bracket, or double quote. If the Authentication Type option is set to use authentication, this password is the authentication password and must include at least 8 characters.
 - Select **SNMPv3 User**.
 - **SNMP User Type.** Select User Access to enable the user to view the SNMP MIB, or Trap Target to enable the user to receive SNMP trap notifications. If you select Trap Target you must set the Trap Host Address option.
 - **Authentication Type.** Select whether to use MD5 or SHA) authentication, or no authentication (None). The default is MD5. Authentication uses the user password.
 - **Privacy Type.** Select whether to use DES or AES encryption, or no encryption (none). The default is none. To use encryption you must also set the Privacy Password and Authentication Type options.
 - **Privacy Password.** If the Privacy Type option is set to use encryption, specify an encryption password, which is case sensitive; can include a maximum of 32 bytes using characters except a backslash, comma, angle bracket, or double quote; and must include at least 8 characters.
 - **Trap Host Address.** If you set the user type to Trap Target, specify the IP address of the host system that will receive SNMP traps.

3. Click **Add User**.

Modifying users

You can change settings either for a general user that can access the WBI (RAIDar), CLI, FTP, or SMI-S interfaces, or for an SNMPv3 user.

To modify a general user

1. In the Configuration View panel, right-click the system and select **Configuration > Users > Modify User**. A table displays details for each user. For each interface a user can access, a checkmark appears in the WBI, CLI, SNMP, FTP, and SMI-S columns.
2. In the main panel, select the user to modify.
3. Set the options:
 - **Password.** A password is case sensitive and can include a maximum of 32 bytes using characters except a backslash, comma, angle bracket, or double quote.
 - **User Roles.** Select Monitor to let the user view system settings, or Manage to let the user view and change system settings. You cannot change the roles of user manage.
 - **User Type.** Select an option to identify the user's experience level: Standard, Advanced, or Diagnostic. This parameter does not affect access to commands.
 - **WBI Access.** Allows access to the web-based management interface.
 - **CLI Access.** Allows access to the command-line management interface.

- **FTP Access.** Allows access to the FTP interface, which can be used instead of the WBI to install firmware updates and download logs.
- **SMI-S Access.** Allows access to the SMI-S interface, used for management of the system through your network.
- **Base Preference.** Select the base for entry and display of storage-space sizes, either Base 10 or Base 2. In base 2, sizes are shown as powers of 2, using 1024 as a divisor for each magnitude. In base 10, sizes are shown as powers of 10, using 1000 as a divisor for each magnitude. 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.
- **Precision Preference.** Select the number of decimal places (1–10) for display of storage-space sizes.
- **Unit Preference.** Select a unit for display of storage-space sizes: Auto, TB, GB, MB. The Auto option lets the system 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, precision is set to 1, and base is set to 10, the size 0.11709 TB is shown as 117.1 GB. The default is Auto.
- **Temperature Preference.** Specifies the scale to use for temperature values: Celsius or Fahrenheit.
- **Auto Sign Out (minutes).** Select the amount of time that the user's session can be idle before the user is automatically signed out (2–720 minutes). The default is 30 minutes.
- **Locale.** The user's preferred display language, which overrides the system's default display language. Installed language sets include Chinese-Simplified, Chinese-Traditional, Dutch, English, French, German, Italian, Japanese, Korean, and Spanish.

4. Click **Modify User**.

User changes take effect when the user next signs in.

To modify an SNMPv3 user

1. In the Configuration View panel, right-click the system and select **Configuration > Users > Modify User**. A table displays details for each user. SNMP3 users can only access the SNMP interface; the other columns are not applicable.
2. In the main panel, select the user to modify.
3. Set the options:
 - **Password.** A password is case sensitive and can include a maximum of 32 bytes using characters except a backslash, comma, angle bracket, or double quote. If the Authentication Type option is set to use authentication, this password is the authentication password and must include at least 8 characters.
 - **SNMP User Type.** Select User Access to enable the user to view the SNMP MIB, or Trap Target to enable the user to receive SNMP trap notifications. If you select Trap Target you must set the Trap Host Address option.
 - **Authentication Type.** Select whether to use MD5 or SHA authentication, or no authentication (None). The default is MD5. Authentication uses the user password.
 - **Privacy Type.** Select whether to use DES or AES encryption, or no encryption (none). The default is none. To use encryption you must also set the Privacy Password and Authentication Type options.
 - **Privacy Password.** If the Privacy Type option is set to use encryption, specify an encryption password, which is case sensitive; can include a maximum of 32 bytes using characters except a backslash, comma, angle bracket, or double quote; and must include at least 8 characters.
 - **Trap Host Address.** If you set the user type to Trap Target, specify the IP address of the host system that will receive SNMP traps.

4. Click **Modify User**.

User changes take effect when the user next signs in.

Removing users

To remove a user

1. In the Configuration View panel, right-click the system and select **Configuration > Users > Remove User**.
2. In the main panel, select the user to remove. You cannot remove the manage user.
3. Click **Remove User**. A confirmation dialog appears.
4. Click **Remove** to continue; otherwise, click **Cancel**. If you clicked Remove, a processing dialog appears. When processing is complete, the user is removed from the table.
5. Click **OK**.

Configuring system settings

Changing the system date and time

You can enter values manually for the system date and time, or you can set the system to use NTP as explained in ["About the system date and time" \(page 30\)](#).

To use manual date and time settings

1. In the Configuration View panel, right-click the system and select **Configuration > System Settings > Date, Time**. The date and time options appear.
2. Set the options:
 - Time. Enter the time in the format *hh:mm:ss*, where *hh* is the hour (0–23), *mm* is the minutes (0–59), and *ss* is the seconds (0–59).
 - Month. Select the month.
 - Day. Enter the day number.
 - Year. Enter the year using four digits.
 - Network Time Protocol (NTP). Select **Disabled**.
3. Click **Apply**.

To obtain the date and time from an NTP server

1. In the Configuration View panel, right-click the system and select **Configuration > System Settings > Date, Time**. The date and time options appear.
2. Set the options:
 - Network Time Protocol (NTP). Select **Enabled**.
 - NTP Time Zone Offset. Optional. The system's time zone as an offset in hours (-12 through +14) and optionally minutes (0–59) from Coordinated Universal Time (UTC). To specify a positive offset, the '+' is optional; to specify a negative offset, the '-' is required. The hour value can have one or two digits and can omit a leading zero. If the minutes value is specified it must have two digits; if it is omitted, the minutes value is set to 00.
 - NTP Server Address. Optional. If the system should retrieve time values from a specific NTP server, enter the address of an NTP server. If no IP server address is set, the system listens for time messages sent by an NTP server in broadcast mode.
3. Click **Apply**.

Changing host interface settings

To enable the system to communicate with hosts or with remote systems having FC or iSCSI interfaces, you must configure host-interface options. For SAS ports there are no host-interface options.

AssuredSAN 4004 Series CNC controller module host ports can be configured as FC or iSCSI ports. For both controller modules, all ports can be set to FC, all ports can be set to iSCSI, or the first two ports in each controller module can be set to FC and the second two ports can be set to iSCSI. FC ports support use of 16-Gbit/s or 8-Gbit/s SFPs. A 16-Gbit/s SFP can run at 16 Gbit/s, 8 Gb/s, or 4 Gb/s or auto-negotiate its link speed. An 8-Gbit/s SFP can run at 8 Gb/s or 4 Gb/s or auto-negotiate its link speed.

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

To change FC host interface settings

1. In the Configuration View panel, right-click the system and select **Configuration > System Settings > Host Interfaces**.
2. Set the Speed option to the proper value to communicate with the host. The speed can be set to auto (the default), which auto-negotiates the proper link speed with the host, or to 4Gb, 8Gb, or 16Gb (Gbit/s). 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.
3. Click **Apply**.

To change iSCSI host interface settings

1. In the Configuration View panel, right-click the system and select **Configuration > System Settings > Host Interfaces**.
2. In the Common Settings for iSCSI section of the panel, set the options that apply to all iSCSI ports:
 - Authentication (CHAP). Enables or disables use of Challenge Handshake Authentication Protocol. Disabled by default.

NOTE: CHAP records for iSCSI login authentication must be defined if CHAP is enabled. To create CHAP records, see "[Configuring CHAP](#)" (page 76).

- Jumbo Frames. Enables or disables support for jumbo frames. A normal frame can contain 1500 bytes whereas a jumbo frame can contain a maximum of 9000 bytes for larger data transfers. Disabled by default.

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

- iSCSI IP Version. Specifies whether IP values use Internet Protocol version 4 (IPv4) or version 6 (IPv6) format. IPv4 uses 32-bit addresses. IPv6 uses 128-bit addresses. The default is IPv4.
 - iSNS. Enables or disables registration with a specified Internet Storage Name Service server, which provides name-to-IP-address mapping. Disabled by default.
 - iSNS Address. Specifies the IP address of an iSNS server. The default address is all zeroes.
 - Alternate iSNS Address. Specifies the IP address of an alternate iSNS server, which can be on a different subnet. The default address is all zeroes.
3. In the upper section of the panel, set the port-specific options:
 - 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 CNC 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
 - Netmask. For IPv4, subnet mask for assigned port IP address. The default is 255.255.255.0.
 - Gateway. For IPv4, gateway IP address for assigned port IP address. The default is 0.0.0.0.
 - Default Router. For IPv6, default router for assigned port IP address. If the gateway was set for IPv4 and then ports were switched to IPv6, the default is `:: IPv4-address`; otherwise, the default is `::` (the short form of all zeroes).
 - Link-Local Address. For IPv6, the link-local address that is automatically generated from the MAC address and assigned to the port.

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

4. Click **Apply**.

Changing network interface settings

You can configure addressing parameters for each controller's network port. You can set static IP values or use DHCP.

In DHCP mode, network port IP address, subnet mask, and gateway values are obtained 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:

- DHCP: disabled
- 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, the following initial values are set and remain set until the system is able to contact a DHCP server for new addresses:

- Controller IP addresses: 169.254.x.x (where the value of x.x is the lowest 16 bits of the controller serial number)
- IP subnet mask: 255.255.0.0
- Gateway IP address: 0.0.0.0

169.254.x.x addresses (including gateway 169.254.0.1) are on a private subnet that is reserved for unconfigured systems and the addresses are not routable. This prevents the DHCP server from reassigning the addresses and possibly causing a conflict where two controllers have the same IP address. As soon as possible, change these IP values to proper values for your network.

△ **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 Configuration View panel, right-click the system and select **Configuration > System Settings > Network Interfaces**.
2. Set the IP address source to **DHCP**.
3. Click **Apply**. If the controllers successfully obtain IP values from the DHCP server, the new IP values are displayed.
4. Record the new addresses.
5. Sign out and access RAIDar 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 controller.
2. In the Configuration View panel, right-click the system and select **Configuration > System Settings > Network Interfaces**.
3. Set the IP address source to **manual**.
4. Set the values for each controller. You must set a unique IP address for each network port.
5. Record the IP values you assign.
6. Click **Apply**.
7. Sign out and access RAIDar using the new IP addresses.

Setting system information

To set system information

1. In the Configuration View panel, right-click the system and select **Configuration > System Settings > System Information**.
2. In the main panel, set the System Name, System Contact person, System Location, and System Information (description). Each value can include a maximum of 79 bytes, using characters except double quote or backslash. 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.
3. Click **Apply**.

Configuring advanced settings

Changing disk settings

Configuring SMART

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

To change the SMART setting

1. In the Configuration View panel, right-click the system and select **Configuration > Advanced Settings > Disk**.
2. Set the SMART Configuration option to one of the following:
 - **Don't Modify**. Allows current disks to retain their individual SMART settings and does not change the setting for new disks added to the system.
 - **Enabled**. Enables SMART for all current disks after the next rescan and automatically enables SMART for new disks added to the system. This option is the default.
 - **Disabled**. Disables SMART for all current disks after the next rescan and automatically disables SMART for new disks added to the system.
3. Click **Apply**.

Configuring dynamic spares

The dynamic spares feature lets you use all of your disks in fault-tolerant vdisks without designating a disk as a spare. With dynamic spares enabled, if a disk fails and you replace it with a compatible disk, the storage system rescans the bus, finds the new disk, automatically designates it a spare, and starts reconstructing the vdisk. A compatible disk has enough capacity to replace the failed disk and is the same type (SAS SSD, enterprise SAS, or midline SAS). If a dedicated spare, global spare, or compatible available disk is already present, the dynamic spares feature uses that disk to start the reconstruction and the replacement disk can be used for another purpose.

To change the dynamic spares setting

1. In the Configuration View panel, right-click the system and select **Configuration > Advanced Settings > Disk**.
2. Either select (enable) or clear (disable) the **Dynamic Spare Capability** option.
3. Click **Apply**.

Configuring drive spin down for available disks and global spares

The drive spin down (DSD) feature monitors disk activity within system enclosures and spins down inactive disks to conserve energy. You can enable or disable DSD for available disks and global spares, and set the period of inactivity after which available disks and global spares automatically spin down.

To configure a time period to suspend and resume DSD for all disks, see ["Scheduling drive spin down for all disks" \(page 51\)](#). To configure DSD for a vdisk, see ["Configuring drive spin down for a vdisk" \(page 57\)](#).

DSD affects disk operations as follows:

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

To configure DSD for available disks and global spares

1. In the Configuration View panel, right-click the local system and select **Configuration > Advanced Settings > Disk**.
2. Set the options:
 - Either select (enable) or clear (disable) the **Available and Spare Drive Spin Down Capability** option. If you are enabling DSD, a warning prompt appears; to use DSD, click **Yes**; to leave DSD disabled, click **No**.
 - Set the Drive Spin Down Delay (minutes) option, which is the period of inactivity after which available disks and global spares automatically spin down, from 1–360 minutes. If DSD is enabled and no delay value is set, the default is 15 minutes. The value 0 disables DSD.
3. Click **Apply**. When processing is complete a success dialog appears.
4. Click **OK**.

Scheduling drive spin down for all disks

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

To configure DSD for a vdisk, see ["Configuring drive spin down for a vdisk" \(page 57\)](#). To configure DSD for available disks and global spares, see ["Configuring drive spin down for available disks and global spares" \(page 50\)](#).

DSD affects disk operations as follows:

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

To schedule DSD for all disks

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

Configuring the EMP polling rate

You can change the frequency interval at which the storage system polls each attached enclosure's EMP for status changes. Typically you can use the default setting.

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

To change the EMP polling rate

1. In the Configuration View panel, right-click the system and select **Configuration > Advanced Settings > Disk**.
2. Set the EMP Polling Rate interval. The options are 5, 10, or 30 seconds; or 1, 5, 10, 15, 20, 25, 30, 45, or 60 minutes. The default is 5 seconds.
3. Click **Apply**.

Changing system cache settings

Changing the synchronize-cache mode

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

To change the synchronize-cache mode

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

Changing the missing LUN response

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

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

To change the missing LUN response

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

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

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

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

To change host access to the write-back cache setting

1. In the Configuration View panel, right-click the system and select **Configuration > Advanced Settings > Cache**.
2. Either select (enable) or clear (disable) the **Host Control of Write-Back Cache** option.
3. Click **Apply**.

Changing the controllers' cache redundancy mode

In a dual-controller system's default redundancy/operating mode, Active-Active ULP, data for volumes configured to use write-back cache is automatically mirrored between the two controllers. Cache redundancy has a slight impact on performance but provides fault tolerance. You can disable cache redundancy, which permits independent cache operation for each controller; this is called independent cache performance mode (ICPM).

The advantage of ICPM is that the two controllers can achieve very high write bandwidth and still use write-back caching. User data is still safely stored in nonvolatile RAM, with backup power provided by super-capacitors should a power failure occur. This feature is useful for high-performance applications that do not require a fault-tolerant

environment for operation; that is, where speed is more important than the possibility of data loss due to a drive fault prior to a write completion.

The disadvantage of ICPM is that if a controller fails, the other controller will not be able to fail over (that is, take over I/O processing for the failed controller). If a controller experiences a complete hardware failure, and needs to be replaced, then user data in its write-back cache will be lost.

△ **CAUTION:** Data might be compromised if a RAID controller failure occurs after it has accepted write data, but before that data has reached the disk drives. Do not use ICPM in an environment that requires fault tolerance.

NOTE: Enabling ICPM will disable partner firmware upgrade.

To change the controllers' cache redundancy mode

1. In the Configuration View panel, right-click the system and select **Configuration > Advanced Settings > Cache**.
2. Either select (enable) or clear (disable) the **Independent Cache Performance Mode** option. In a dual-controller system this option is disabled by default. In Single Controller mode this option is grayed out.
3. Click **Apply**. The controllers are automatically restarted to change their mode.

Changing auto-write-through cache triggers and behaviors

You can set conditions that cause ("trigger") a controller to change the cache mode from write-back to write-through, as described in ["About volume cache options" \(page 22\)](#). You can also specify actions for the system to take when write-through caching is triggered.

To change auto-write-through cache triggers and behaviors

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

Configuring partner firmware update

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

To change the partner firmware update setting

1. In the Configuration View panel, right-click the system and select **Configuration > Advanced Settings > Firmware**.
2. Either select (enable) or clear (disable) the **Partner Firmware Update** option.
3. Click **Apply**.

Configuring system utilities


Configuring background scrub for vdisks

You can enable or disable whether the system continuously analyzes disks in vdisks to find and fix disk errors. This command will fix parity mismatches for RAID 3, 5, 6, and 50; mirror mismatches for RAID 1 and 10; and media errors for all RAID levels.

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

When a scrub is complete, event 207 is logged and specifies whether errors were found and whether user action is required.

Enabling background vdisk scrub is recommended.

 **TIP:** If you choose to disable background vdisk scrub, you can still scrub a selected vdisk by using Tools > Media Scrub Vdisk ([page 87](#)).

To configure background scrub for vdisks

1. In the Configuration View panel, right-click the system and select **Configuration > Advanced Settings > System Utilities**.
2. Set the options:
 - Either select (enable) or clear (disable) the **Vdisk Scrub** option. This option is disabled by default.
 - Set the Vdisk Scrub Interval (hours) option, which is the interval between background vdisk scrub finishing and starting again, from 1–360 hours; the default is 24 hours.
3. Click **Apply**.

Configuring background scrub for disks not in vdisks

You can enable or disable whether the system continuously analyzes disks that are not in vdisks to find and fix disk errors. Background disk scrub will start 72 hours after it is enabled. The interval between background disk scrub finishing and starting again is 72 hours.

Enabling background disk scrub is recommended.

To configure background scrub for disks not in vdisks

1. In the Configuration View panel, right-click the system and select **Configuration > Advanced Settings > System Utilities**.
2. Either select (enable) or clear (disable) the **Disk Scrub** option. This option is disabled by default.
3. Click **Apply**.

Configuring utility priority

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

To change the utility priority

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

Enabling/disabling managed logs

You can enable or disable the managed logs feature, which allows log files to be transferred from the storage system to a log-collection system to avoid losing diagnostic data. For an overview of the managed logs feature, see ["About managed logs" \(page 32\)](#). Before enabling log management, configure log destinations as described in ["Configuring email notification" \(page 43\)](#) and ["Configuring SNMP notification" \(page 44\)](#). After enabling log management, you can test it as described in ["Testing notifications" \(page 86\)](#).

To enable or disable managed logs

1. In the Configuration View panel, right-click the system and select **Configuration > Advanced Settings > System Utilities**.
2. Either select (enable) or clear (disable) the **Managed Logs** option. This option is disabled by default.
3. Click **Apply**.

Configuring remote systems

Adding a remote system

You can add a management object to obtain information from a remote storage system. This allows a local system to track remote systems by their network-port IP addresses and cache their login credentials. The IP address can then be used in commands that need to interact with the remote system.

To add a remote system

1. In the Configuration View panel, either:
 - Right-click the local system and select **Configuration > Remote System > Add Remote System**.
 - Right-click a remote system and select **Configuration > Add Remote System**.
2. In the main panel set the options:
 - IP address. IP address of a network port on the remote system.
 - User Name. User name for a user that has Manage-level access on the remote system.
 - Password. Optional. Password for the specified user.
3. Click **Create Remote System**. If the task succeeds, the new remote system appears in the Configuration View panel.

Deleting remote systems

You can delete the management objects for remote systems.

After establishing replication to a remote system, if you choose to delete the remote system you can safely do so without affecting replications. However, because the remote system's name and IP address will no longer appear in user interfaces, record this information before deleting the remote system so that you can access it at a later time, such as to delete old replication images or for disaster recovery.

To delete remote systems

1. In the Configuration View panel, either:
 - Right-click the local system and select **Configuration > Remote System > Delete Remote System**.
 - Right-click a remote system and select **Configuration > Delete Remote System**.
2. In the main panel, select the remote systems to remove. To select or clear all remote systems, toggle the checkbox in the heading row.
3. Click **Delete Remote System(s)**. A confirmation dialog appears.
4. Click **Delete** to continue; otherwise, click **Cancel**. If you clicked Delete, a processing dialog appears. If the task succeeds, the System Overview panel and a success dialog appear.
5. Click **OK**. As processing completes, the deleted items are removed from the Configuration View panel.

Configuring a vdisk

Managing dedicated spares

You can assign a maximum of four available disks to a fault-tolerant vdisk (RAID 1, 3, 5, 6, 10, 50) for use as spares by that vdisk only. A spare must be the same type (SAS SSD, enterprise SAS, or midline SAS) as other disks in the vdisk, and have sufficient capacity to replace the smallest disk in the vdisk.

NOTE: If you upgraded from an earlier release that did not distinguish between enterprise and midline SAS disks, you might have vdisks that contain both types of disks. For such a vdisk, whose RAID-level label has the suffix **-MIXED** in the Configuration View panel, you can designate either or both types of disks to be spares.

If a disk in the vdisk fails, a dedicated spare is automatically used to reconstruct the vdisk. A fault-tolerant vdisk other than RAID-6 becomes Critical when one disk fails. A RAID-6 vdisk becomes Degraded when one disk fails and Critical when two disks fail. After the vdisk's parity or mirror data is completely written to the spare, the vdisk returns to fault-tolerant status. For RAID-50 vdisks, if more than one sub-vdisk becomes critical, reconstruction and use of assigned spares occur in the order sub-vdisks are numbered.

To change a vdisk's spares

1. In the Configuration View panel, right-click a vdisk and select **Configuration > Manage Dedicated Spares**. The main panel shows information about the selected vdisk, its spares, and all disks in the system. Existing spares are labeled SPARE.
 - In the Disk Sets table, the number of white slots in the Disks column of the SPARE row shows how many spares you can add to the vdisk.
 - In the Graphical or Tabular view, only existing spares and suitable available disks are selectable.
2. Select spares to remove, disks to add as spares, or both. To add a spare, select its checkbox. To remove a spare, clear its checkbox.
3. Click **Modify Spares**. If the task succeeds, the panel is updated to show which disks are now spares for the vdisk.

Changing a vdisk's name

To change a vdisk's name

1. In the Configuration View panel, right-click a vdisk and select **Configuration > Modify Vdisk Name**. The main panel shows the vdisk's name.
2. Enter a new name. A vdisk name is case sensitive; cannot already exist in the system; and cannot include a comma, double quote, angle bracket, or backslash. The name you enter can have a maximum of 17 bytes.
3. Click **Modify Name**. The new name appears in the Configuration View panel.

Changing a vdisk's owner

Each vdisk is owned by one of the controllers, A or B, known as the *preferred owner*. Typically, you should not need to change vdisk ownership.

When a controller fails, the partner controller assumes temporary ownership of the failed controller's vdisks and resources, becoming the *current owner*. If the system uses a fault-tolerant cabling configuration, both controllers' LUNs are accessible through the partner.

△ CAUTION:

- Before changing the owning controller for a vdisk, you must stop host I/O to the vdisk's volumes.
- Because a volume and its snap pool must be in vdisks owned by the same controller, if an ownership change will cause volumes and their snap pools to be owned by different controllers, the volumes will not be able to access their snap pools.

Changing the owner of a vdisk does not affect the mappings volumes in that vdisk.

To change a vdisk's owner

1. In the Configuration View panel, right-click a vdisk and select **Configuration > Modify Vdisk Owner**. The main panel shows the vdisk's owner.
2. Select a new owner.
3. Click **Modify Owner**. A confirmation dialog appears.
4. Click **Yes** to continue; otherwise, click **No**. If you clicked Yes, a processing dialog appears. When processing is complete a success dialog appears.
5. Click **OK**.

Configuring drive spin down for a vdisk

The drive spin down (DSD) feature monitors disk activity within system enclosures and spins down inactive disks to conserve energy. For a specific vdisk, you can enable or disable DSD and set the period of inactivity after which the vdisk's disks and dedicated spares automatically spin down.

To configure a time period to suspend and resume DSD for all vdisks, see ["Scheduling drive spin down for all disks" \(page 51\)](#). To configure DSD for available disks and global spares, see ["Configuring drive spin down for available disks and global spares" \(page 50\)](#).

DSD affects disk operations as follows:

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

To configure DSD for a vdisk

1. In the Configuration View panel, right-click a vdisk and select **Configuration > Configure Vdisk Drive Spin Down**.
2. Set the options:
 - Either select (enable) or clear (disable) the **Enable Drive Spin Down** option.
 - Set the Drive Spin Down Delay (minutes) option, which is the period of inactivity after which the vdisk's disks and dedicated spares automatically spin down, from 1–360 minutes.
3. Click **Apply**. When processing is complete a success dialog appears.
4. Click **OK**.

Configuring a volume

Changing a volume's name

To change a volume's name

1. In the Configuration View panel, right-click a volume and select **Configuration > Modify Volume Name**.
2. Enter a new name. A volume name is case sensitive; cannot already exist in a vdisk; cannot include a comma, double quote, angle bracket or backslash; and can have a maximum of 20 bytes.
3. Click **Modify Name**. The new name appears in the Configuration View panel.

Changing a volume's cache settings

For explanations of volume cache options, see ["About volume cache options" \(page 22\)](#).

△ 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.
- 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.

To change a volume's cache settings

1. In the Configuration View panel, right-click a volume and select **Configuration > Modify Volume Cache Settings**.
2. In the main panel, set the read-ahead cache options:
 - Write Policy. Select Write-back or Write-through. The default is Write-back.
 - Write Optimization. Select Standard or No-mirror. Atomic write is not supported in this release. The default is Standard.
 - Read Ahead Size. Select Disabled, Adaptive, Stripe, or a specific size (512KB, 1MB, 2MB, 4MB, 8MB, 16MB, or 32MB).
3. Click **Modify Cache Settings**.

Configuring a snapshot

Changing a snapshot's name

To change a snapshot's name

1. In the Configuration View panel, right-click a snapshot and select **Configuration > Modify Snapshot Name**.
2. Enter a new name. A snapshot name is case sensitive; cannot already exist in a vdisk; cannot include a comma, double quote, or backslash; and can have a maximum of 20 bytes.
3. Click **Modify Name**. The new name appears in the Configuration View panel.

Configuring a snap pool

Changing a snap pool's name


To change a snap pool's name

1. In the Configuration View panel, right-click a snap pool and select **Configuration > Modify Snap Pool Name**.
2. Enter a new name. A snap pool name is case sensitive; cannot already exist in a vdisk; cannot include a comma, double quote, or backslash; and can have a maximum of 20 bytes.
3. Click **Modify Name**. The new name appears in the Configuration View panel.

3 Provisioning the system

Using the Provisioning Wizard

The Provisioning Wizard helps you create a vdisk with volumes and to map the volumes to hosts. Before using this wizard, read documentation for your product to learn about vdisks, volumes, and mapping. Then plan the vdisks and volumes you want to create and the default mapping settings you want to use.

The wizard guides you through the following steps. For each step you can view help by clicking the help icon  in the wizard panel. As you complete steps they are highlighted at the bottom of the panel. If you cancel the wizard at any point, no changes are made.

- Specify a name and RAID level for the vdisk
- Select disks to use in the vdisk
- Specify the number and size of volumes to create in the vdisk
- Specify the default mapping for access to the volume by hosts
- Confirm changes and apply them

Step 1: Starting the wizard

1. In the Configuration View panel, right-click the system and select either **Provisioning > Provisioning Wizard** or **Wizards > Provisioning Wizard**. The wizard panel appears.
2. Click **Next** to continue.

Step 2: Specifying the vdisk name and RAID level

A *vdisk* is a virtual disk that is composed of one or more disks, and has the combined capacity of those disks. The number of disks that a vdisk can contain is determined by its RAID level. When you create a vdisk, all its disks must be the same type: either SAS SSD, enterprise SAS, or midline SAS. A maximum of 16 vdisks per controller can exist.

A vdisk can contain different models of disks, and disks with different capacities. If you mix disks with different capacities, the smallest disk determines the logical capacity of all other disks in the vdisk, regardless of RAID level. For example, the capacity of a vdisk composed of one 500-GB disk and one 750-GB disk is equivalent to a vdisk composed of two 500-GB disks. To maximize capacity, use disks of similar size. For greatest reliability, use disks of the same size and rotational speed.

In a single-controller system, all vdisks are owned by that controller. In a dual-controller system, when a vdisk is created the system automatically assigns the owner to balance the number of vdisks each controller owns; or, you can select the owner. Typically it doesn't matter which controller owns a vdisk.

In a dual-controller system, when a controller fails, the partner controller assumes temporary ownership of the failed controller's vdisks and resources. If the system uses a fault-tolerant cabling configuration, both controllers' LUNs are accessible through the partner.

When you create a vdisk you can also create volumes within it. A volume is a logical subdivision of a vdisk, and can be mapped to controller host ports for access by hosts. The storage system presents only volumes, not vdisks, to hosts.

To create a vdisk

1. Set the options:
 - Vdisk name. Optionally change the default name for the vdisk. A vdisk name is case sensitive; cannot already exist in the system; and cannot include a comma, double quote, angle bracket or backslash. The name you enter can have a maximum of 17 bytes.
 - Assign to. If the system is operating in Active-Active ULP mode, optionally select a controller to be the preferred owner for the vdisk. The default, Auto, automatically assigns the owner to load-balance vdisks between controllers. If the system is operating in Single Controller mode, the Assign to setting is ignored and the system automatically load-balances vdisks in anticipation of the insertion of a second controller in the future.
 - RAID level. Select a RAID level for the vdisk. Each RAID level requires a minimum number of disks, so the number of available disks determines which RAID-level options are selectable. The default, with at least three available disks, is RAID 5.

- Number of sub-vdisks. For a RAID-10 or RAID-50 vdisk, optionally change the number of sub-vdisks that the vdisk should contain.
- Chunk size. For RAID 5, 6, 10, or 50, optionally set the amount of contiguous data that is written to a vdisk member before moving to the next member of the vdisk. For RAID 50, this option sets the chunk size of each RAID-5 sub-vdisk. The chunk size of the RAID-50 vdisk is calculated as: $configured_chunk_size \times (subvdisk_members - 1)$. For RAID 1, chunk size has no meaning and is therefore disabled. The default size is 512KB.

2. Click **Next** to continue.

Step 3: Selecting disks

Select disks to include in the vdisk. The Disk Selection Sets table has one row for each sub-vdisk in a RAID-10 or RAID-50 vdisk, or a single row for a vdisk having another RAID level. The table also has a SPARE row where you can assign dedicated spares to the vdisk. In each row, the Disks field shows how many disks you can, and have, assigned. As you select disks, the table shows the amount of storage space in the vdisk. For descriptions of storage-space color codes, see "[About storage-space color codes](#)" (page 31).

The Tabular tab shows all available disks in all enclosures in a table, displaying Health, Name, Type, State, Size, Enclosure, Serial Number, and Status. The Graphical tab shows disk information graphically, displaying the state for each disk (VDISK, AVAIL, SPARE). Only available disks can be selected. Disks you select are highlighted and color-coded to match the rows in the Disk Selection Sets table. Based on the type of disk you select first (SAS SSD, enterprise SAS, or midline SAS), only available disks of that type become selectable; you cannot mix disk types when creating a vdisk.

To select disks and spares

1. Select disks to populate each vdisk row. When you have selected enough disks, a checkmark appears in the table's Complete field.
2. Optionally select up to four dedicated spares for the vdisk.
3. Click **Next** to continue.

Step 4: Defining volumes

A *volume* is a logical subdivision of a vdisk and can be mapped to controller host ports for access by hosts. A mapped volume provides the storage for a file system partition you create with your operating system or third-party tools. The storage system presents only volumes, not vdisks, to hosts.

You can create multiple volumes with the same base name, size, and default mapping settings. If you choose to define volumes in this step, you will define their mapping settings in the next step.

To define volumes

1. Set the options:
 - Number of volumes to create. Specify the number of volumes to create. If you do not want to create volumes, enter **0**. After changing the value, press **Tab**.
 - Volume size. Specify the size of each volume. The default size is the total capacity of the vdisk divided by the number of volumes.
 - Base name for volumes. Specify the base name for the volumes. A volume name is case sensitive; cannot already exist in a vdisk; cannot include a comma, double quote, angle bracket, or backslash; and can have a maximum of 20 bytes.
2. Click **Next** to continue.

Step 5: Setting the default mapping

Specify *default mapping* settings to control whether and how hosts will be able to access the vdisk's volumes. These settings include:

- A logical unit number (LUN), used to identify a mapped volume to hosts. Both controllers share one set of LUNs. Each LUN can be assigned as the default LUN for only one volume in the storage system; for example, if LUN 5 is the default for Volume 1, LUN5 cannot be the default LUN for any other volume.

- The level of access — read-write, read-only, or no access — that hosts will have to each volume. When a mapping specifies no access, the volume is *masked*.
- Controller host ports through which hosts will be able to access each volume. To maximize performance, it is recommended to map a volume to at least one host port on the controller that the volume's vdisk is assigned to. To sustain I/O in the event of controller failure, it is recommended to map to at least one host port on each controller.

After a volume is created you can change its default mapping, and create, modify, or delete explicit mappings. An *explicit mapping* overrides the volume's default mapping for a specific host.

NOTE: When mapping a volume to a host using the Linux ext3 file system, specify read-write access; otherwise, the file system will be unable to mount/present/map the volume and will report an error such as "unknown partition table."

To specify the default mapping

1. Select **Map**.
2. Set the starting LUN for the volumes. If this LUN is available, it will be assigned to the first volume and the next available LUNs in sequence will be assigned to any remaining volumes.
3. In the enclosure view or list, select controller host ports through which attached hosts can access each volume.
4. Select the access level that hosts will have to each volume: **read-write**, **read-only**, or **no-access** (masked).
5. Click **Next** to continue.

Step 6: Confirming vdisk settings

Confirm that the values 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 setting changes and finish the wizard.

Creating a vdisk

To create a vdisk

1. In the Configuration View panel, right-click the system or **Vdisks** and then select **Provisioning > Create Vdisk**.
2. In the main panel set the options:
 - **Vdisk name.** Optionally change the default name for the vdisk. A vdisk name is case sensitive; cannot already exist in the system; and cannot include a comma, double quote, angle bracket, or backslash. The name you enter can have a maximum of 17 bytes.
 - **Assign to.** If the system is operating in Active-Active ULP mode, optionally select a controller to be the preferred owner for the vdisk. The default, Auto, automatically assigns the owner to load-balance vdisks between controllers. If the system is operating in Single Controller mode, the Assign to setting is ignored and the system automatically load-balances vdisks in anticipation of the insertion of a second controller in the future.
 - **RAID Level.** Select a RAID level for the vdisk. Each RAID level requires a minimum number of disks, so the number of available disks determines which RAID-level options are selectable. The default, with at least three available disks, is RAID 5.
 - **Number of Sub-vdisks.** For a RAID-10 or RAID-50 vdisk, optionally change the number of sub-vdisks that the vdisk should contain.
 - **Chunk size.** For RAID 5, 6, 10, or 50, optionally set the amount of contiguous data that is written to a vdisk member before moving to the next member of the vdisk. For RAID 50, this option sets the chunk size of each RAID-5 sub-vdisk. The chunk size of the RAID-50 vdisk is calculated as: $configured_chunk_size \times (subvdisk_members - 1)$. For RAID 1, chunk size has no meaning and is therefore disabled. The default size is 512KB.
 - **Online Initialization.** If you select (enable) this option, you can use the vdisk while it is initializing but because the verify method is used to initialize the vdisk, initialization takes more time. If you clear (disable) this option, you must wait for initialization to complete before using the vdisk, but initialization takes less time. Online initialization is fault tolerant.

3. Select disks to include in the vdisk. The Disk Selection Sets table has one row for each sub-vgdisk in a RAID-10 or RAID-50 vdisk, or a single row for a vdisk having another RAID level. The table also has a SPARE row where you can assign dedicated spares to the vdisk. In each row, the Disks field shows how many disks you can, and have, assigned. As you select disks, the table shows the amount of storage space in the vdisk. For descriptions of storage-space color codes, see ["About storage-space color codes" \(page 31\)](#). The Tabular tab shows all available disks in all enclosures in a table, displaying Health, Name, Type, State, Size, Enclosure, Serial Number, and Status. The Graphical tab shows disk information graphically, displaying the state for each disk (VDISK, AVAIL, SPARE). Only available disks can be selected. Disks you select are highlighted and color-coded to match the rows in the Disk Selection Sets table. Based on the type of disk you select first (SAS SSD, enterprise SAS, or midline SAS), only available disks of that type become selectable; disks of different types cannot be mixed in a vdisk.

To select disks and spares:

- Select disks to populate each vdisk row. When you have selected enough disks, a checkmark appears in the table's Complete field.
 - Optionally select up to four dedicated spares for the vdisk.
4. Click **Create Vdisk**. If the task succeeds, the new vdisk appears in the Configuration View panel.

Deleting vdisks

△ **CAUTION:** Deleting a vdisk removes all of its volumes and their data.

To delete vdisks

1. Verify that hosts are not accessing volumes in the vdisks that you want to delete.
2. In the Configuration View panel, either:
 - Right-click the system or **Vdisks** and then select **Provisioning > Delete Vdisks**.
 - Right-click a vdisk and select **Provisioning > Delete Vdisk**.
3. In the main panel, select the vdisks to delete. To select or clear all vdisks, toggle the checkbox in the heading row.
4. Click **Delete Vdisk(s)**. A confirmation dialog appears.
5. Click **Delete** to continue; otherwise, click **Cancel**. If you clicked Delete, a processing dialog appears. If the task succeeds, an overview panel and a success dialog appear.
6. Click **OK**. As processing completes, the deleted items are removed from the Configuration View panel.

Managing global spares

You can designate a maximum of 16 global spares for the system. If a disk in any fault-tolerant vdisk (RAID 1, 3, 5, 6, 10, 50) fails, a global spare is automatically used to reconstruct the vdisk. At least one vdisk must exist before you can add a global spare. A spare must have sufficient capacity to replace the smallest disk in an existing vdisk.

The vdisk remains in critical status until the parity or mirror data is completely written to the spare, at which time the vdisk returns to fault-tolerant status. For RAID-50 vdisks, if more than one sub-vgdisk becomes critical, reconstruction and use of spares occur in the order sub-vgdisks are numbered.

To change the system's global spares

1. In the Configuration View panel, right-click the system and select **Provisioning > Manage Global Spares**. The main panel shows information about available disks in the system. Existing spares are labeled GLOBAL SP.
 - In the Disk Sets table, the number of white slots in the Disks field shows how many spares you can add.
 - In the Graphical or Tabular view, only existing global spares and suitable available disks are selectable.
2. Select spares to remove, disks to add as spares, or both.
3. Click **Modify Spares**. If the task succeeds, the panel is updated to show which disks are now global spares.

Creating a volume set

In a vdisk that has sufficient free space, you can create multiple volumes with the same base name and size. Optionally, you can specify a default mapping for the volumes; otherwise, they will be created unmapped.

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

To create a volume set

1. In the Configuration View panel, right-click a vdisk and select **Provisioning > Create Volume Set**.
2. In the main panel, set the options:
 - Volume Set Base-name. Optionally change the default base name for the volumes. The volume names will consist of the base name and a number that increments from 000. If a name in the series is already in use, the next name in the series is assigned. For example, for a two-volume set starting with Volume000, if Volume001 already exists, the second volume is named Volume002. A base name is case sensitive; cannot already be used by another vdisk; cannot include a comma, double quote, angle bracket, or backslash; and can have a maximum of 16 bytes.
 - Total Volumes. Specify the number of volumes to create. Volumes are created up to the maximum number supported per vdisk.
 - Size. Optionally change the volume size. The default size is the total space divided by the number of volumes.
 - Map. Select this option to specify a default mapping for the volumes:
 - Access. Select the access level that hosts will have to the volumes.
 - LUN. If the access level is set to read-write or read-only, set a LUN for the first volume. The next available LUN is assigned to the next volume mapped through the same ports. If a LUN to be assigned to a volume is already in use, that volume and any subsequent volumes are not mapped.
 - In the enclosure view or list, select controller host ports through which attached hosts can access the volumes.
3. Click **Apply**. If the task succeeds, the new volumes appear in the Configuration View panel.

Creating a volume

You can add a volume to a vdisk that has sufficient free space, and define default mapping settings.

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

NOTE: In rare cases, a large amount of I/O can cause a snap pool that is too small to fill quickly. This can result in all snapshots being deleted due to the snap pool running out of space. Create snap pools of at least 50 GB to avoid this situation.

To create a volume in a vdisk

1. In the Configuration View panel, right-click a vdisk and select **Provisioning > Create Volume**.
2. In the main panel, set the options:
 - Volume name. Optionally change the default name. A volume name is case sensitive; cannot already exist in a vdisk; cannot include a comma, double quote, or backslash; and can have a maximum of 20 bytes.
 - Size. Optionally change the default size, which is all free space in the vdisk.
 - Enable Snapshots. If the system is licensed to use Snapshots and you want to create snapshots of this volume, select this option. This specifies to create the volume as a master volume instead of as a standard volume, and enables the Snap Pool and Replication Prepare options.

- Snap Pool. Select either:
 - Standard Policy. This option creates a snap pool named `spvolume-name` whose size is either 20% of the volume size or 6 GB, whichever is larger.
 - Reserve Size. Specify the size of the snap pool to create in the vdisk and associate with the new volume. The default size is either 20% of the volume size or 6 GB, whichever is larger.
 - Attach Pool. Select an existing snap pool to associate with the new volume.
 - Replication Prepare. If the system is licensed to use remote replication and you want to use this volume as a secondary volume, select this option. Selecting this option disables the Map option.
 - Map. Select this option to change the default mapping for the volume:
 - Access. Select the access level that hosts will have to the volume.
 - LUN. If the access level is set to read-write or read-only, set a LUN for the volume.
 - In the enclosure view or list, select controller host ports through which attached hosts can access the volume.
3. Click **Apply**. If the task succeeds, the new volume appears in the Configuration View panel. If you specified an option to create a snap pool, the new snap pool also appears in that panel.

Deleting volumes

You can use the Delete Volumes panel to delete standard and master volumes.

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

To delete volumes

1. Verify that hosts are not accessing the volumes that you want to delete.
2. In the Configuration View panel, either:
 - Right-click the system or **Vdisks** or a vdisk and then select **Provisioning > Delete Volumes**.
 - Right-click a volume and select **Provisioning > Delete Volume**.
3. In the main panel, select the volumes to delete. To select up to 100 volumes or clear all selections, toggle the checkbox in the heading row.
4. Click **Delete Volume(s)**.
5. Click **Delete** to continue; otherwise, click **Cancel**. If you clicked Delete, a processing dialog appears. If the task succeeds, an overview panel and a success dialog appear.
6. Click **OK**. As processing completes, the deleted items are removed from the Configuration View panel.

NOTE: The system might be unable to delete a large number of volumes in a single operation. If you specified to delete a large number of volumes, verify that all were deleted. If some of the specified volumes remain, repeat the deletion on those volumes.

Changing default mapping for multiple volumes

For all volumes in all vdisks or a selected vdisk, you can change the default access to those volumes by all hosts. When multiple volumes are selected, LUN values are sequentially assigned starting with a LUN value that you specify. For example, if the starting LUN value is 1 for 30 selected volumes, the first volume's mapping is assigned LUN 1 and so forth, and the last volume's mapping is assigned LUN 30. For LUN assignment to succeed, ensure that no value in the sequence is already in use.

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

NOTE: You cannot map the secondary volume of a replication set.

NOTE: When mapping a volume to a host using the Linux ext3 file system, specify read-write access; otherwise, the file system will be unable to mount/present/map the volume and will report an error such as “unknown partition table.”

To change default mapping for multiple volumes

1. In the Configuration View panel, right-click **Vdisks** or a vdisk and then select **Provisioning > Map Volume Defaults**. In the main panel, a table shows all the volumes for the selected vdisk.
2. In the table, select the volumes to change. To select up to 100 volumes or clear all selections, toggle the checkbox in the heading row.
3. Either:
 - Map the volumes to all hosts by setting a starting LUN, selecting ports, and setting access to **read-only** or **read-write**.
 - Mask the volumes from all hosts by setting a starting LUN, selecting ports, and setting access to **no-access**. Setting the default mapping to **no-access** will result in the LUN mapping being removed.
4. Click **Apply**. A message specifies whether the change succeeded or failed.
5. Click **OK**.

Explicitly mapping multiple volumes

For all volumes in all vdisks or a selected vdisk, you can change access to those volumes by a specific host. When multiple volumes are selected, LUN values are sequentially assigned starting with a LUN value that you specify. For example, if the starting LUN value is 1 for 30 selected volumes, the first volume’s mapping is assigned LUN 1 and so forth, and the last volume’s mapping is assigned LUN 30. For LUN assignment to succeed, ensure that no value in the sequence is already in use. When specifying access through specific ports, the ports and host must be the same type (for example, FC).

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

NOTE: You cannot map the secondary volume of a replication set.

NOTE: When mapping a volume to a host using the Linux ext3 file system, specify read-write access; otherwise, the file system will be unable to mount/present/map the volume and will report an error such as “unknown partition table.”

To explicitly map multiple volumes

1. In the Configuration View panel, right-click **Vdisks** or a vdisk and then select **Provisioning > Map Volumes**. In the main panel, a table shows all the volumes for the selected vdisk.
2. In the table, select the volumes to change. To select up to 100 volumes or clear all selections, toggle the checkbox in the heading row.
3. In the Hosts table, select the host to change access for.

4. Either:
 - Map the volumes to the host by setting a starting LUN, selecting ports, and setting access to **read-only** or **read-write**.
 - Mask the volumes from the host by setting a starting LUN, selecting ports, and setting access to **no-access**.
5. Click **Apply**. A message specifies whether the change succeeded or failed.
6. Click **OK**.

Changing a volume's default mapping

△ **CAUTION:** Volume mapping changes take effect immediately. Make changes that limit access to volumes when the volumes are not in use. Be sure to unmount/unpresent/unmap a volume before changing the volume's LUN.

NOTE: You cannot map the secondary volume of a replication set.

NOTE: When mapping a volume to a host using the Linux ext3 file system, specify read-write access; otherwise, the file system will be unable to mount/present/map the volume and will report an error such as "unknown partition table."

To view the default mapping

In the Configuration View panel, right-click a volume and select **Provisioning > Default Mapping**. The main panel shows the volume's default mapping:

- Ports. Controller host ports through which the volume is mapped to the host.
- LUN. Volume identifier presented to the host.
- Access. Volume access type: read-write, read-only, not-mapped.

To modify the default mapping

1. Select **Map**.
2. Set the LUN and select the ports and access type. Setting the default mapping to **no-access** will result in the LUN mapping being removed.
3. Click **Apply**. A message specifies whether the change succeeded or failed.
4. Click **OK**. Each mapping that uses the default settings is updated.

To delete the default mapping

1. Clear **Map**.
2. Click **Apply**. A message specifies whether the change succeeded or failed.
3. Click **OK**. Each mapping that uses the default settings is updated.

Changing a volume's explicit mappings

△ **CAUTION:** Volume mapping changes take effect immediately. Make changes that limit access to volumes when the volumes are not in use. Be sure to unmount/unpresent/unmap a volume before changing the volume's LUN.

NOTE: You cannot map the secondary volume of a replication set.

NOTE: When mapping a volume to a host using the Linux ext3 file system, specify read-write access; otherwise, the file system will be unable to mount/present/map the volume and will report an error such as "unknown partition table."

To view volume mappings

In the Configuration View panel, right-click a volume and select **Provisioning > Explicit Mappings**. The main panel shows the following information about the volume's mappings:

- Type. Explicit or Default. Settings for an explicit mapping override the default mapping.
- Host ID. WWPN or IQN.
- Host Name. User-defined nickname for the host.
- Ports. Controller host ports through which the host is mapped to the volume.
- LUN. Volume identifier presented to the host.
- Access. Volume access type: read-write, read-only, no-access (masked), or not-mapped.

To create an explicit mapping

1. In the Maps for Volume table, select a host.
2. Select **Map**.
3. Set the LUN and select the ports and access type.
4. Click **Apply**. A message specifies whether the change succeeded or failed.
5. Click **OK**. The mapping becomes Explicit with the new settings.

To modify an explicit mapping

1. In the Maps for Volume table, select the Explicit mapping to change.
2. Set the LUN and select the ports and access type.
3. Click **Apply**. A message specifies whether the change succeeded or failed.
4. Click **OK**. The mapping settings are updated.

To delete an explicit mapping

1. In the Maps for Volume table, select the Explicit mapping to delete.
2. Clear **Map**.
3. Click **Apply**. A message specifies whether the change succeeded or failed.
4. Click **OK**. The mapping returns to the Default mapping.

Unmapping volumes

You can delete all of the default and explicit mappings for multiple volumes.

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

To unmap volumes

1. In the Configuration View panel, right-click **Vdisks** or a vdisk and then select **Provisioning > Unmap Volumes**. In the main panel, a table shows all the volumes for the selected vdisk.
2. In the table, select the volumes to unmap. To select up to 100 items or clear all selections, toggle the checkbox in the heading row.
3. Click **Unmap Volume(s)**. A message specifies whether the change succeeded or failed.
4. Click **OK**. Default and explicit mappings are deleted and the volumes' access type changes to not-mapped.

Expanding a volume

You can expand a standard volume if its vdisk has free space and sufficient resources. Because volume expansion does not require I/O to be stopped, the volume can continue to be used during expansion.

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

NOTE: This command is not supported for master volumes.

To expand a volume

1. In the Configuration View panel, right-click a standard volume and select **Tools > Expand Volume**.
2. In the main panel, specify the amount of free space to add to the volume.
3. Click **Expand Volume**. If the specified value exceeds the amount of free space in the vdisk, a dialog lets you expand the volume to the limit of free space in the vdisk. If the task succeeds, the volume's size is updated in the Configuration View panel.

Creating multiple snapshots

If the system is licensed to use Snapshots, you can select multiple volumes and immediately create a snapshot of each volume.

The first time a snapshot is created of a standard volume, the volume is converted to a master volume and a snap pool is created in the volume's vdisk. The snap pool's size is either 20% of the volume size or 6 GB, whichever is larger. Before creating or scheduling snapshots, verify that the vdisk has enough free space to contain the snap pool.

NOTE: In rare cases, a large amount of I/O can cause a snap pool that is too small to fill quickly. This can result in all snapshots being deleted due to the snap pool running out of space. If you expect the snap pool to have a high rate of data change, view the snap pool's size. If it is less than 50 GB, expand the snap pool to at least 50 GB.

To create multiple snapshots

1. In the Configuration View panel, right-click the system or **Vdisks** or a vdisk and then select **Provisioning > Create Multiple Snapshots**.
2. In the main panel, select each volume to take a snapshot of. To select up to 100 volumes or clear all selections, toggle the checkbox in the heading row.
3. Click **Create Snapshots**. If the task succeeds, the snapshots appear in the Configuration View panel.

Creating a snapshot

If the system is licensed to use Snapshots, you can create a snapshot now or schedule the snapshot task.

The first time a snapshot is created of a standard volume, the volume is converted to a master volume and a snap pool is created in the volume's vdisk. The snap pool's size is either 20% of the volume size or 6 GB, whichever is larger. Before creating or scheduling snapshots, verify that the vdisk has enough free space to contain the snap pool.

NOTE: In rare cases, a large amount of I/O can cause a snap pool that is too small to fill quickly. This can result in all snapshots being deleted due to the snap pool running out of space. If you expect the snap pool to have a high rate of data change, view the snap pool's size. If it is less than 50 GB, expand the snap pool to at least 50 GB.

To create a snapshot now

1. In the Configuration View panel, right-click a volume and select **Provisioning > Create Snapshot**.
2. In the main panel, select **Now**.
3. Optionally change the default name for the snapshot. A snapshot name is case sensitive; cannot already exist in a vdisk; cannot include a comma, double quote, or backslash; and can have a maximum of 20 bytes.
4. Click **Create Snapshot**. If the task succeeds, the snapshot appears in the Configuration View panel.

To schedule a create snapshot task

1. In the Configuration View panel, right-click a volume and select **Provisioning > Create Snapshot**.
2. In the main panel, select **Scheduled**.
3. Set the options:
 - Snapshot prefix. Optionally change the default prefix to identify snapshots created by this task. The prefix is case sensitive; cannot include a comma, double quote, or backslash; and can have a maximum of 14 bytes. Automatically created snapshots are named *prefix_sn*, where *n* starts at 001.
 - Snapshots to Retain. Select the number of snapshots to retain. 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 volume's oldest snapshot is unmapped, reset, and renamed to the next name in the sequence.
 - Start Schedule. 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.
 - Date must use the format *yyyy-mm-dd*.
 - 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.
 - Recurrence. Specify the interval at which the task should run. Set the interval to at least two minutes. For better performance if this task will run under heavy I/O conditions or on more than three volumes, set the retention count and the schedule interval to similar values; for example if the retention count is 10 then the interval should be set to 10 minutes. The default is 1 minute.
 - Time Constraint. Specify a time range within which the task should run.
 - Date Constraint. Specify days when the task should run. Ensure that this constraint includes the Start Schedule date.
 - End Schedule. Specify when the task should stop running.
4. Click **Schedule Snapshots**. If processing succeeds, the schedule is saved and can be viewed in the overview panel for the volume or system.

Deleting snapshots

You can use the Delete Snapshots panel to delete standard and replication snapshots.

When you delete a snapshot, all data uniquely associated with that snapshot is deleted and associated space in the snap pool is freed for use. Snapshots can be deleted in any order, irrespective of the order in which they were created.

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

△ **CAUTION:** If a replication snapshot's type is shown as a "sync point" for its replication set, consider carefully whether you want to delete that snapshot. If you delete the current sync point, then if a replication-set failure occurs, a prior sync point will be used. If you delete the only sync point then the next replication will require a full sync to be performed (*all* data to be re-replicated from the primary volume to a secondary volume).

To delete snapshots

1. Verify that hosts are not accessing the snapshots that you want to delete.
2. In the Configuration View panel, right-click either the system or a vdisk or a master volume or a primary volume or a secondary volume or a snapshot or a replication image and then select **Provisioning > Delete Snapshot**.
3. In the main panel, select the snapshots to delete.
4. Click **Delete Snapshot(s)**.
5. Click **OK** to continue; otherwise, click **Cancel**. If you clicked OK, a processing dialog appears. If the task succeeds, an overview panel and a success dialog appear.
6. Click **OK**. As processing completes, the deleted items are removed from the Configuration View panel.

Resetting a snapshot

If the system is licensed to use Snapshots, 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's name and mapping settings are not changed. The snapshot data is stored in the source volume's snap pool. This task is not allowed for a replication snapshot.

△ **CAUTION:** To avoid data corruption, before resetting a snapshot it must be unmounted/unpresented/unmapped from hosts.

You can reset a snapshot now or schedule the reset task.

To reset a snapshot now

1. Unmount/unpresent/unmap the snapshot from hosts.
2. In the Configuration View panel, right-click a snapshot and select **Provisioning > Reset Snapshot**.
3. In the main panel, select **Now**.
4. Click **Reset Snapshot**. A confirmation dialog appears.
5. Click **Yes** to continue; otherwise, click **No**. If you clicked Yes, a processing dialog appears. When processing is complete a success dialog appears.
6. Click **OK**.
7. Optionally, remount/re-present/remap the snapshot.

To schedule a reset snapshot task

1. In the Configuration View panel, right-click a snapshot and select **Provisioning > Reset Snapshot**.
2. In the main panel, select **Scheduled**.
3. Set the options:
 - **Start Schedule.** 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.
 - Date must use the format *yyyy-mm-dd*.
 - 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.
 - **Recurrence.** Specify the interval at which the task should run. Set the interval to at least 2 minutes. The default is 1 minute.
 - **Time Constraint.** Specify a time range within which the task should run.
 - **Date Constraint.** Specify days when the task should run. Ensure that this constraint includes the Start Schedule date.
 - **End Schedule.** Specify when the task should stop running.
4. Click **Reset Snapshot**. If the task succeeded, the schedule is saved and can be viewed in the overview panel for the snapshot or system.
5. Make a reminder to unmount/unpresent/unmap the snapshot before the scheduled task runs.

Creating a volume copy

If the system is licensed to use Volume Copy, you can copy a volume or a snapshot to a new standard volume. The destination volume must be in a vdisk owned by the same controller as the source volume. If the source volume is a snapshot, you can choose whether to include its modified data (data written to the snapshot since it was created). The destination volume is completely independent of the source volume.

The first time a volume copy is created of a standard volume, the volume is converted to a master volume and a snap pool is created in the volume's vdisk. The snap pool's size is either 20% of the volume size or 6 GB, whichever is larger. Before creating or scheduling copies, verify that the vdisk has enough free space to contain the snap pool.

NOTE: In rare cases, a large amount of I/O can cause a snap pool that is too small to fill quickly. This can result in all snapshots being deleted due to the snap pool running out of space. If you expect the snap pool to have a high rate of data change, view the snap pool's size. If it is less than 50 GB, expand the snap pool to at least 50 GB.

For a master volume, the volume copy creates a transient snapshot, copies the data from the snapshot, and deletes the snapshot when the copy is complete. For a snapshot, the volume copy is performed directly from the source; this source data may change if modified data is to be included in the copy and the snapshot is mounted/presented/mapped and I/O is occurring to it.

To ensure the integrity of a copy of a master volume, unmount/unpresent/unmap the volume or at minimum perform a system cache flush and refrain from writing to the volume. Since the system cache flush is not natively supported on all operating systems, it is recommended to unmount/unpresent/unmap temporarily. The volume copy is for all data on the disk at the time of the request, so if there is data in the operating-system cache, that will not be copied over. Unmounting/unpresenting/unmapping the volume forces the cache flush from the operating system. After the volume copy has started, it is safe to remount/re-present/remap the volume and/or resume I/O.

To ensure the integrity of a copy of a snapshot with modified data, unmount/unpresent/unmap the snapshot or perform a system cache flush. The snapshot will not be available for read or write access until the volume copy is complete, at which time you can remount/re-present/remap the snapshot. If modified write data is not to be included in the copy, then you may safely leave the snapshot mounted/presented. During a volume copy using snapshot modified data, the system takes the snapshot offline, as shown by the Snapshot Overview panel.

The volume copy's progress is shown in the Volume Overview panel.

You can create a volume copy now or schedule the copy task.

To create a volume copy now

1. In the Configuration View panel, right-click a volume and select **Provisioning > Create Volume Copy**.
2. In the main panel, select **Now**.
3. Set the options:
 - **New Volume Name.** Optionally change the default name for the destination volume. A volume name is case sensitive; cannot already exist in a vdisk; cannot include a comma, double quote, or backslash; and can have a maximum of 20 bytes.
 - **Residing On Vdisk.** Optionally change the destination vdisk.
 - **With Modified Data.** If the source volume is a snapshot, select this option to include the snapshot's modified data in the copy. Otherwise, the copy will contain only the data that existed when the snapshot was created.
4. Click **Copy the Volume**. A confirmation dialog appears.
5. Click **Yes** to continue; otherwise, click **No**. If you clicked Yes and With Modified Data is selected and the snapshot has modified data, a second confirmation dialog appears.
6. Click **Yes** to continue; otherwise, click No. If you clicked Yes, the volume copy operation starts. While the operation is in progress, the destination volume is offline and its type is shown as "standard*". If you unmounted/unpresented/unmapped a snapshot to copy its modified data, *wait* until processing is complete before you remount/re-present/remap it. If the task succeeds, the destination volume's type becomes standard and the volume appears in the Configuration View panel.
7. Optionally map the volume to hosts.

To schedule a volume copy task

1. In the Configuration View panel, right-click a volume and select **Provisioning > Create Volume Copy**.
2. In the main panel, select **Scheduled**.
3. Set the options:
 - **New Volume Prefix.** Optionally change the default prefix to identify volumes created by this task. The prefix is case sensitive; cannot include a comma, double quote, or backslash; and can have a maximum of 14 bytes. Automatically created volumes are named *prefix_cn*, where *n* starts at 001.
 - **Residing On Vdisk.** Optionally change the destination vdisk.
 - **With Modified Data.** If the source volume is a snapshot, select this option to include the snapshot's modified data in the copy. Otherwise, the copy will contain only the data that existed when the snapshot was created.
 - **Start Schedule.** 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.
 - Date must use the format *yyyy-mm-dd*.
 - 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.
 - **Recurrence.** Specify interval at which the the task should run. Set the interval to at least 2 minutes. The default is 1 minute.
 - **Time Constraint.** Specify a time range within which the task should run.
 - **Date Constraint.** Specify days when the task should run. Ensure that this constraint includes the Start Schedule date.
 - **End Schedule.** Specify when the task should stop running.
4. Click **Schedule Volume Copy**. If the task succeeded, the schedule is saved and can be viewed in the overview panel for the volume or system.
5. If you will copy snapshot modified data, make a reminder to unmount/unpresent/unmap the snapshot before the scheduled task runs.

Aborting a volume copy

If the system is licensed to use Volume Copy, you can cancel an in-progress volume copy operation. When the cancellation is complete, the destination volume is deleted.

To abort a volume copy

1. In the Configuration View panel, right-click the destination volume and then select **Provisioning > Abort Volume Copy**. The Volume Overview panel shows the operation's progress.
2. Click **Abort Volume Copy**. A message confirms that the operation has been aborted.
3. Click **OK**. The destination volume is removed from the Configuration View panel.

Rolling back a volume

You can roll back (revert) the data in a volume to the data that existed when a specified snapshot was created. You also have the option of including its modified data (data written to the snapshot since it was created). For example, you might want to take a snapshot, mount/present/map it for read/write, and then install new software on the snapshot for testing. If the software installation is successful, you can roll back the volume to the contents of the modified snapshot.

△ CAUTION:

- Before rolling back a volume you must unmount/unpresent/unmap it from data hosts to avoid data corruption. If you want to include snapshot modified data in the roll back, you must also unmount/unpresent/unmap the snapshot.
- If the snap pool runs out of space, the master volume will change to read only until the rollback has completed.
- Whenever you perform a roll back, the data that existed on the volume is replaced by the data on the snapshot; that is, all data on the volume written since the snapshot was taken is lost. As a precaution, take a snapshot of the volume before starting a roll back.

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

During a roll back operation using snapshot modified data, the snapshot must be unmounted/unpresented/unmapped and cannot be accessed. Unmounting/unpresenting/unmapping the snapshot ensures that all data cached by the host is written to the snapshot; if unmounting/unpresenting/unmapping is not performed at the host level prior to starting the roll back, data may remain in host cache, and thus not be rolled back to the master volume. As a precaution against inadvertently accessing the snapshot, the system also takes the snapshot offline, as shown by the Snapshot Overview panel. The snapshot becomes inaccessible in order to prevent any data corruption to the master volume. The snapshot can be remounted/re-presented/remapped once the roll back is complete.

To roll back a volume

1. Unmount/unpresent/unmap the volume from hosts.
2. If the roll back will include snapshot modified data, unmount/unpresent/unmap the snapshot from hosts.
3. In the Configuration View panel, right-click a volume and select **Provisioning > Roll Back Volume**.
4. In the main panel, set the options:
 - For Volume.
 - From Snapshot Volume. Enter the name of the snapshot to roll back to.
 - **With Modified Data**. Select this option to include the snapshot's modified data in the roll back. Otherwise, the master volume will contain only the data that existed when the snapshot was created.
5. Click **Roll Back Volume**. The roll back starts. You can now remount/re-present/remap the volume.
6. When the roll back is complete, if you unmounted/unpresented/unmapped the snapshot you can remount/re-present/remap it.

Creating a snap pool

Before you can convert a standard volume to a master volume or create a master volume for snapshots, a snap pool must exist. A snap pool and its associated master volumes can be in different vdisks, but must be owned by the same controller.

To create a snap pool

1. In the Configuration View panel, right-click a vdisk and select **Provisioning > Create Snap Pool**.
2. In the main panel set the options:
 - Snap Pool name. Optionally change the default name for the snap pool. A snap pool name is case sensitive; cannot already exist in the system; and can have a maximum of 20 bytes. A name cannot include a comma, double quote, or backslash.
 - Size. Optionally change the default size, which is all free space in the vdisk. The minimum size is 5.37 GB.

NOTE: In rare cases, a large amount of I/O can cause a snap pool that is too small to fill quickly. This can result in all snapshots being deleted due to the snap pool running out of space. Create snap pools of at least 50 GB to avoid this situation.

3. Click **Create Snap Pool** If the task succeeds, the new snap pool appears in the Configuration View panel.

Deleting snap pools

Before you can delete a snap pool you must delete any associated snapshots, and either delete the associated master volume or convert the master volume to a standard volume.

To delete snap pools

1. Verify that no master volume or snapshots are associated with the snap pool.
2. In the Configuration View panel, either:
 - Right-click the local system or **Vdisks** or a vdisk and select **Provisioning > Delete Snap Pools**.
 - Right-click a snap pool and select **Provisioning > Delete Snap Pool**.
3. In the main panel, select the snap pools to delete.
4. Click **Delete Snap Pool(s)**.
5. Click **Delete** to continue; otherwise, click **Cancel**. If you clicked Delete, a processing dialog appears. If the task succeeds, an overview panel and a success dialog appear.
6. Click **OK**. As processing completes, the deleted items are removed from the Configuration View panel.

Adding a host

To add a host

1. Determine the host's WWPN or IQN.
2. In the Configuration View panel, right-click the system or **Hosts** and then select **Provisioning > Add Host**.
3. In the main panel set the options:
 - Host ID (WWN/IQN). Enter the host's WWPN or IQN. A WWPN value can include a colon between each pair of digits but the colons will be discarded.
 - Host Name. Optionally change the default name to one that helps you easily identify the host; for example, FileServer_1. A host name is case sensitive; cannot already exist in the system; cannot include a comma, double quote, angle bracket, or backslash; and can have a maximum of 15 bytes.
4. Click **Add Host**. If the task succeeds, the new host appears in the Configuration View panel.

Removing hosts

To remove hosts

1. Verify that the hosts you want to remove are not accessing volumes.
2. In the Configuration View panel, either:
 - Right-click the system or **Hosts** and then select **Provisioning > Remove Hosts**.
 - Right-click a host and select **Provisioning > Remove Host**.
3. In the main panel, select the hosts to remove. To select or clear all items, toggle the checkbox in the heading row.
4. Click **Remove Host(s)**. A confirmation dialog appears.
5. Click **Remove** to continue; otherwise, click **Cancel**. If you clicked Remove, a processing dialog appears. If the task succeeds, an overview panel and a success dialog appear.
6. Click **OK**. As processing completes, the deleted items are removed from the Configuration View panel.

Changing a host's name

To change a host's name

1. In the Configuration View panel, right-click a host and select **Provisioning > Rename Host**.
2. Enter a new name that helps you easily identify the host; for example, FileServer_1. A host name is case sensitive; cannot already exist in the system; cannot include a comma, double quote, or backslash; and can have a maximum of 15 bytes.
3. Click **Modify Name**.

Changing host mappings

For each volume that is mapped to the selected host, you can create, modify, and delete explicit mappings. To change a volume's default mapping, see ["Changing a volume's default mapping" \(page 66\)](#).

△ **CAUTION:** Volume mapping changes take effect immediately. Make changes that limit access to volumes when the volumes are not in use. Be sure to unmount/unpresent/unmap a volume before changing the volume's LUN.

NOTE: You cannot map the secondary volume of a replication set.

NOTE: When mapping a volume to a host using the Linux ext3 file system, specify read-write access; otherwise, the file system will be unable to mount/present/map the volume and will report an error such as "unknown partition table."

To view host mappings

In the Configuration View panel, right-click a host and select **Provisioning > Manage Host Mappings**. The main panel shows the following information about volumes mapped to the host:

- Type. Explicit or Default. Settings for an explicit mapping override the default mapping.
- Name. Volume name.
- Serial Number. Volume serial number.
- Ports. Controller host ports through which the volume is mapped to the host.
- LUN. Volume identifier presented to the host.
- Access. Volume access type: read-write, read-only, no-access (masked), or not-mapped.

To create an explicit mapping

1. In the Maps for Host table, select the Default mapping to override.
2. Select **Map**.
3. Set the LUN and select the ports and access type.
4. Click **Apply**. A message specifies whether the change succeeded or failed.
5. Click **OK**. The mapping becomes Explicit with the new settings.

To modify an explicit mapping

1. In the Maps for Host table, select the Explicit mapping to change.
2. Set the LUN and select the ports and access type.
3. Click **Apply**. A message specifies whether the change succeeded or failed.
4. Click **OK**. The mapping settings are updated.

To delete an explicit mapping

1. In the Maps for Host table, select the Explicit mapping to delete.
2. Clear **Map**.
3. Click **Apply**. A message specifies whether the change succeeded or failed.
4. Click **OK**. The mapping returns to the Default mapping.

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 entries must exist on each device. Each CHAP entry 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 storage system, the host is the initiator and the storage system is the target.

To enable or disable CHAP for all iSCSI hosts, see ["Changing host interface settings" \(page 47\)](#).

To add or modify a CHAP entry

1. In the Configuration View panel, right-click **Hosts** or a specific host and then select **Provisioning > Configure CHAP**. If any CHAP entries exist, a table shows them by node name.
2. Optionally, select an entry whose name you want to change to create a new entry. The entry's values appear in the option fields.
3. Set the options:
 - Node Name (IQN). The initiator name, typically in IQN format.
 - Secret. The secret that the target uses to authenticate the initiator. The secret is case sensitive and can include 12–16 bytes.
 - Name, if mutual CHAP. Optional; for mutual CHAP only. Specifies the target name, typically in IQN format. The value is case sensitive, can include a maximum of 223 bytes, and must differ from the initiator name. To find a controller iSCSI port's IQN, select the controller enclosure, view the Enclosure Overview panel ([page 110](#)), select the Rear Graphical tab, select an iSCSI port, and view the Target ID field.
 - Secret, if mutual CHAP. Optional; for mutual CHAP only. Specifies the secret that the initiator uses to authenticate the target. The secret is case sensitive, can include 12–16 bytes, and must differ from the initiator secret. A storage system's secret is shared by both controllers.
4. Click **Add/Modify Entry**. If the task succeeds, the new or modified entry appears in the CHAP entries table.

To delete a CHAP entry

1. In the Configuration View panel, right-click **Hosts** or a specific host and then select **Provisioning > Configure CHAP**. If any CHAP entries exist, a table shows them by node name.
2. Select the entry to delete.
3. Click **Delete Entry**. If the task succeeds, the entry is removed from the CHAP entries table.

Modifying a schedule

To modify a schedule

1. In the Configuration View panel, right-click the system or a volume or a snapshot and select **Provisioning > Modify Schedule**. In the main panel, a table shows each schedule.
2. In the table, select the schedule to modify. For information about schedule status values, see "[Schedule properties](#)" (page 105).
3. Set the options:
 - **Snapshot Prefix.** Optionally change the default prefix to identify snapshots created by this task. The prefix is case sensitive; cannot include a comma, double quote, or backslash; and can have a maximum of 14 bytes. Automatically created snapshots are named *prefix_sn*, where *n* starts at 001.
 - **Snapshots to Retain.** Select the number of snapshots to retain. 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 volume's oldest snapshot is unmapped, reset, and renamed to the next name in the sequence.
 - **Start Schedule.** 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.
 - Date must use the format *yyyy-mm-dd*.
 - 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.
 - **Recurrence.** Specify interval at which the task should run. Select either One Time or how often the task should occur. If the task is recurrent, select Minutes, Hours, Days, Weeks, Months, or Years from the list. If One Time is selected, no further options are available and the task will occur only on the date and time specified in Start Schedule.
 - For a snapshot schedule, set the interval to at least 2 minutes. For better performance if this task will run under heavy I/O conditions or on more than three volumes, set the retention count and the interval to similar values; for example if the retention count is 10 then set the interval to 10 minutes.
 - For a volume-copy or reset-snapshot schedule, set the interval to at least 2 minutes.
 - For a replication schedule, set the interval to at least 30 minutes.
 - **Time Constraint.** Specify a time range within which the task should run. Select either No Time Constraint or times between which the task will run.
 - If No Time Constraint is selected, the task will run whenever scheduled.
 - If a set of times is specified, the task can only occur during that period of time.
 - **Date Constraint.** Specify days when the task should run. Ensure that this constraint includes the Start Schedule date. Select No Date Constraint, Any, or a specific time.
 - If No Date Constraint is selected, the task will occur whenever scheduled.
 - For Any, select a type of day (any day, weekday, weekend day, or specific day of the week) and a year, month, or specific month. For example, if you select Any Weekday of June, the task can occur only on weekdays in June.
 - For a specific time, select a type of day (any day, weekday, weekend day, or specific day of the week), a number, and a year, month, or specific month. For example, if you select Sunday number 3 of January, the task can occur only on the 3rd Sunday of January.
 - **End Schedule.** Specify when the task should stop running. Select Continuous, Date, or End After.
 - If Continuous is selected, the task will never end.
 - If a date and time is specified, the task will not run as scheduled after that date.
 - If End After is selected, the task will end after the running the number of times you specify. For example, if you enter 5, the task will run only 5 times.
4. Click **Modify Schedule**.
5. Click **Yes** to continue; otherwise, click **No**. If you clicked Yes, a processing dialog appears. When processing is complete a success dialog appears.
6. Click **OK**.

Deleting schedules

If a component has a scheduled task that you no longer want to occur, you can delete the schedule. When a component is deleted, its schedules are also deleted.

To delete task schedules


1. In the Configuration View panel, right-click the system or a volume or a snapshot and select **Provisioning > Delete Schedule**.
2. In the main panel, select the schedule to remove.
3. Click **Delete Schedule**. A confirmation dialog appears.
4. Click **Yes** to continue; otherwise, click **No**. If you clicked Yes, a processing dialog appears. If the task succeeds, the schedules are removed from the table and from the Configuration View panel. When processing is complete a success dialog appears.
5. Click **OK**.

4 Using system tools

Updating firmware

You can view the current versions of firmware in controller modules, expansion modules, and disks, and install new versions.

To monitor the progress of a firmware-update operation by using the activity progress interface, see ["Using the activity progress interface"](#) (page 81) below.

 **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 vdisk is quarantined, resolve the problem that is causing the vdisk to be quarantined before updating firmware. See information about events 172 and 485 in the Event Descriptions Reference Guide, and ["Removing a vdisk from quarantine"](#) (page 88).
 - 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 93](#)).
-

Updating controller-module firmware

A controller enclosure can contain one or two controller modules. In a dual-controller system, both controllers should run the same firmware version. Storage systems in a replication set must run the same firmware version. You can update the firmware in each controller module by loading a firmware file obtained from the enclosure vendor.

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

For best results, the storage system should be 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. If the storage system has a single controller, stop I/O to vdisks before starting the firmware update.
3. Restart the Management Controller (MC) in the controller to be updated; or if PFU is enabled, restart the MCs in both controllers. For the procedure, see ["Restarting or shutting down controllers"](#) (page 85).
4. In the Configuration View panel, right-click the system and select **Tools > Update Firmware**. The table titled Current Controller Versions shows the currently installed versions.
5. Click **Browse** and select the firmware file to install.
6. Click **Install Controller-Module Firmware File**. A dialog box 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 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.

If the Storage Controller cannot be updated, the update operation is cancelled. 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 Management Controller will restart. Until the restart is complete, the RAIDar Sign In page will say that the system is currently unavailable. When this message is cleared, you may sign in.

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

7. Clear your web browser's cache, then sign in to RAIDar. 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: After firmware update has completed on both controllers, if the system health is Degraded and the health reason indicates that the firmware version is incorrect, verify that you specified the correct firmware file and repeat the update. If this problem persists, contact technical support.

Updating expansion-module firmware

A drive enclosure can contain one or two expansion modules. Each expansion module contains an enclosure management processor (EMP). All modules of the same model should run the same firmware version.

Expansion-module firmware is updated in either of two ways:

- When you update controller-module firmware, all expansion modules are automatically updated to a compatible firmware version.
- You can update the firmware in each expansion module by loading a firmware file obtained from the enclosure vendor.

📌 **IMPORTANT:** Disable PFU before attempting to update the firmware. If PFU is not disabled, it will downgrade the firmware on the corresponding expansion module. If this occurs, restart each SC.

To update expansion-module firmware

1. Obtain the appropriate firmware file and download it to your computer or network.
2. If the storage system has a single controller, stop I/O to vdisks before starting the firmware update.
3. In the Configuration View panel, right-click the system and select **Tools > Update Firmware**. The table titled Current Versions of All Expansion Modules (EMPs) shows the currently installed versions.
4. Select the expansion modules to update.
5. Click **Browse** and select the firmware file to install.
6. Click **Install Expansion-Module Firmware File**. Messages show firmware-update progress.

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

It typically takes 2.5 minutes to update each EMP in a drive enclosure. Wait for a message that the code load has completed.

7. Verify that each updated expansion module has the correct firmware version.

Updating disk firmware

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

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

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

To update disk firmware

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. 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.
4. In the Configuration View panel, right-click the system and select **Tools > Update Firmware**. The table titled Current Versions (Revisions) of All Disk Drives shows the currently installed versions.
5. Select the disks to update.
6. Click **Browse** and select the firmware file to install.
7. Click **Install Disk Firmware File**. A dialog box shows firmware-update progress.

△ **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 completed.

8. If the updated disks must be power cycled:
 - a. Shut down both controllers; see ["Restarting or shutting down controllers"](#) (page 85).
 - b. Power cycle all enclosures as described in your product's Setup Guide.
9. Verify that each disk has the correct firmware revision.

Using the activity progress interface

The activity progress interface reports whether a firmware update operation is active and shows the progress through each step of the operation. When the update operation completes, status is presented that either indicates successful completion or indicates an error if the operation failed.

When the interface is enabled it is accessible at all times, except when an MC is restarting. The interface is accessed through an HTTP-based query to a specified MC or to both MCs. The content is presented in a unified way so that activity in both MCs in a dual-controller system can be accessed simultaneously.

Use of this interface will not interfere in any way with firmware update performed via RAIDar or FTP.

To access the activity progress interface

1. Enable the Activity Progress Monitor service; see ["Changing management interface settings"](#) (page 42).
2. In a new tab in your web browser, enter a URL of the form:

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

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

Activity progress output

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

Table 11 Activity progress properties and values

Property	Value
Time	The date and time of the latest status update.
Seconds	The number of seconds this component has been active.
Component	The name of the object being processed.
Status	The status of the component representing its progress/completion state. <ul style="list-style-type: none">• ACTIVE: The operation for this component is currently active and in progress.• OK: The operation for this component completed successfully and is now inactive.• N/A: The operation for this component was not completed because it was not applicable.• ERROR: The operation for this component failed with an error (see code and message).
Code	A numeric code indicating the status. <ul style="list-style-type: none">• 0: The operation for this component completed with a "completed successfully" status.• 1: The operation for this component was not attempted because it is not applicable (the component doesn't exist or doesn't need updating).• 2: The operation is in progress. The other properties will indicate the progress item (message, current, total, percent).• 10 or higher: The operation for this component completed with a failure. The code and message indicate the reason for the error.
Message	A textual message indicating the progress status or error condition.

Saving logs

To help service personnel diagnose a system problem, you might be asked to provide system log data. Using RAIDar, you can save log data to a compressed zip file. The file will contain the following data:

- 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

NOTE: The controllers share one memory buffer for gathering log data and for loading firmware. Do not try to perform more than one save-logs operation at a time, or to perform a firmware-update operation while performing a save-logs operation.

To save logs

1. In the Configuration View panel, right-click the system and select **Tools > Save Logs**.
2. In the main panel:
 - a. Enter your name, email address, and phone number so support personnel will know who provided the log data. These entries can each be 100 bytes long using characters except a backslash, angle bracket, or double quote.
 - b. Enter comments that describe the problem and specify the date and time when the problem occurred. This information helps service personnel when they analyze the log data. Comment text can be 500 bytes long using characters except a backslash, angle bracket, or double quote.
3. Click **Save Logs**.

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.

Log data is collected, which takes several minutes.

4. When prompted to open or save the file, click **Save**.
 - If you are using Firefox and have a download directory set, the file `store.zip` is saved there.
 - 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: Because the file is compressed, you must uncompress it before you can view the files it contains. To examine diagnostic data, first view `store_yyyy_mm_dd_hh_mm_ss.logs`.

Resetting a host port

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 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).

For an FC host port configured to use FC-AL (loop) topology, a reset issues a loop initialization primitive (LIP). For SAS, resetting a host port issues a `COMINIT/COMRESET` sequence and might reset other ports.

To reset a host port

1. In the Configuration View panel, right-click the system and select **Tools > Reset Host Port**.
2. Select the port or port pair to reset.
3. Click **Reset Host Port**.

Rescanning disk channels

A rescan forces a rediscovery of disks and enclosures in the storage system. If both Storage Controllers are online and able to communicate with both expansion modules in each connected enclosure, rescan rebuilds the internal SAS layout information, reassigns enclosure IDs of attached enclosures based on controller A's enclosure cabling order, and ensures that the enclosures are displayed in the proper order. A manual rescan temporarily pauses all I/O processes, then resumes normal operation. It can take up to two minutes for the enclosure IDs to be corrected.

A manual rescan may be needed after system power-up to display enclosures in the proper order. Whenever you replace a drive chassis or controller chassis, perform a manual rescan to force fresh discovery of all drive enclosures connected to the controller enclosure.

A manual rescan is not needed after inserting or removing 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. In the Configuration View panel, right-click the system and select **Tools > Rescan Disk Channels**.
3. Click **Rescan**.

Restoring system defaults

If the system is not working properly and you cannot determine why, you can restore its default configuration settings. You then can reconfigure the settings that are necessary to use the system.

To restore defaults, use the CLI's `restore defaults` command, as described in the CLI Reference Guide.

Clearing disk metadata

△ CAUTION:

- Only use this command when all vdisks are online and leftover disks exist. Improper use of this command may result in data loss.
- Do not use this command when a vdisk is offline and one or more leftover disks exist.
- If you are uncertain whether to use this command, contact technical support for further assistance.

Each disk in a vdisk has metadata that identifies the owning vdisk, the other members of the vdisk, and the last time data was written to the vdisk. The following situations cause a disk to become a *leftover*:

- Vdisk members' 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.

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

- The disk's health becomes Degraded and its How Used state becomes LEFTOVR.
- The disk is automatically excluded from the vdisk, causing the vdisk's health to become Degraded or Fault, depending on the RAID level.
- The disk's fault LED is illuminated amber.

If spares are available, and the health of the vdisk is Degraded, the vdisk 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 How Used state to AVAIL, making the disk available for use in a new vdisk or as a spare.

If spares are not available to begin reconstruction, or reconstruction has not completed, keep the leftover disk so that you'll 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 Configuration View panel, right-click the system and then select **Tools > Clear Disk Metadata**.
2. In the main panel, select leftover disks to clear metadata from. To select or clear all leftover disks, toggle the checkbox in the heading row.
3. Click **Clear Metadata**. When processing is complete a success dialog appears.
4. Click **OK**.

Restarting or shutting down controllers

You can restart the processors in a controller module when RAIDar informs you that you have changed a configuration setting that requires restarting or when the controller is not working properly. Shut down the processors in a controller module before you remove it from an enclosure, or before you power off its enclosure for maintenance, repair, or a move.

A restart can be performed on either the Storage Controller processor or the Management Controller processor. A shut down affects both processors.

Restarting

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

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

△ **CAUTION:** If you restart both Storage Controllers, both Management Controllers will also be restarted so all users will lose access to the system and its data until the restart is complete.

NOTE: When a Storage Controller is restarted, live performance statistics that it recorded will be reset; historical performance statistics are not affected. In a dual-controller system, disk statistics may be reduced but will not be reset to zero, because disk statistics are summed between the two controllers. For more information, see help for commands that show statistics.

To perform a restart

1. In the Configuration View panel, right-click the local system and select **Tools > Shut Down or Restart Controller**.
2. In the main panel, set the options:
 - Operation. Select **Restart**.
 - Controller Type. Select the type of controller processor to restart: **Management** or **Storage**.
 - Controller. Select the controller to restart: **A**, **B**, or **Both**.
3. Click **Restart now**. A confirmation dialog appears.
4. Click **Yes** to continue; otherwise, click **No**. If you clicked Yes, a second confirmation dialog appears.
5. Click **Yes** to continue; otherwise, click **No**. If you clicked Yes, a message describes restart activity.

Shutting down

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 the Storage Controller in both controller modules is shut down, hosts cannot access the system's data. Perform a shut down before removing a controller module or powering down the system.

△ **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. In the Configuration View panel, right-click the local system and select **Tools > Shut Down or Restart Controller**.
2. In the main panel, set the options:
 - Operation. Select **Shut down**.
 - Controller. Select the controller to shut down: **A**, **B**, or **Both**.
3. Click **Shut down now**. A confirmation dialog appears.
4. Click **Yes** to continue; otherwise, click **No**. If you clicked Yes, a second confirmation dialog appears.
5. Click **Yes** to continue; otherwise, click **No**. If you clicked Yes, a message describes shutdown activity.

Testing notifications

You can send test messages to verify that email, SNMP, and/or syslog settings are properly configured for destinations to receive event notifications and managed-logs notifications.

For event notification, the email, SNMP, or syslog settings must include a notification level other than "none (disabled)." For managed-logs notification, the managed logs feature must be configured and enabled. For an overview of the log-management feature, see "[About managed logs](#)" (page 32).

To test event notification

1. In the Configuration View panel, right-click the local system and select **Tools > Test Event Notifications and Managed Logs**.
2. Under the Test Event Notifications heading, click **Send Event**. If the task succeeds, verify that the test message reached the destinations.

To test managed-logs notification

1. In the Configuration View panel, right-click the local system and select **Tools > Test Event Notifications and Managed Logs**.
2. Under the Test Managed Logs Notifications heading, click **Send Managed Logs**. If the task succeeds, verify that the test message reached the destination.

Expanding a vdisk

You can expand the capacity of a vdisk by adding disks to it, up to the maximum number of disks that the storage system supports. Host I/O to the vdisk can continue while the expansion proceeds. You can then create or expand a volume to use the new free space, which becomes available when the expansion is complete. You can expand only one vdisk at a time. The RAID level determines whether the vdisk can be expanded and the maximum number of disks the vdisk can have. This task cannot be performed on an NRAID or RAID-1 vdisk.

Adding single-ported disks to a vdisk that contains dual-ported disks is supported; however, because single-ported disks are not fault-tolerant, a confirmation prompt will appear.

📌 **IMPORTANT:** Expansion can take hours or days to complete, depending on the vdisk's RAID level and size, disk speed, utility priority, and other processes running on the storage system. You can stop expansion only by deleting the vdisk.

Before expanding a vdisk

Back up the vdisk's data so that if you need to stop expansion and delete the vdisk, you can move the data into a new, larger vdisk.

To expand a vdisk

1. In the Configuration View panel, right-click a vdisk and select **Tools > Expand Vdisk**. Information appears about the selected vdisk and all disks in the system.
 - In the Disk Selection Sets table, the number of white slots in the vdisk's Disks field shows how many disks you can add to the vdisk.
 - In the enclosure view or list, only suitable available disks are selectable.
2. Select disks to add.
3. Click **Expand Vdisk**. A processing dialog appears.
4. Click **OK**. The expansion's progress is shown in the **View > Overview** panel.

Verifying a vdisk

If you suspect that a fault-tolerant (mirror or parity) vdisk has a problem, run the Verify utility to check the vdisk's integrity. For example, if the storage system was operating outside the normal temperature range, verify its vdisks. The Verify utility analyzes the selected vdisk to find and fix inconsistencies between its redundancy data and its user data. This utility fixes parity mismatches for RAID 3, 5, 6, and 50, and mirror mismatches for RAID 1 and 10. This task can be performed only on a vdisk whose status is FTOL (fault tolerant and online); it cannot be performed for NRAID or RAID 0.

 **TIP:** Media Scrub Vdisk ([page 87](#)) operates similarly to Verify Vdisk but can find and fix media errors for any RAID level, including NRAID and RAID 0.

Verification can last over an hour, depending on the size of the vdisk, the utility priority, and the amount of I/O activity. You can use a vdisk 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 vdisk is going bad. For information about identifying a failing disk, use the SMART option (see "[Configuring SMART](#)" ([page 50](#))).

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 vdisk

1. In the Configuration View panel, right-click a fault-tolerant vdisk and select **Tools > Verify Vdisk**.
2. Click **Start Verify Utility**. A message confirms that verification has started.
3. Click **OK**. The panel shows the verification's progress.

To abort vdisk verification

1. In the Configuration View panel, right-click a fault-tolerant vdisk and select **Tools > Verify Vdisk**.
2. Click **Abort Verify Utility**. A message confirms that verification has been aborted.
3. Click **OK**.

Scrubbing a vdisk

The system-level Vdisk Scrub option (see "[Configuring background scrub for vdisks](#)" ([page 54](#))) automatically checks all vdisks for disk defects. If this option is disabled, you can still perform a scrub on a selected vdisk. Scrub analyzes a vdisk to find and fix disk errors. It will fix parity mismatches for RAID 3, 5, 6, and 50; 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 vdisk, the utility priority, and the amount of I/O activity. However, a "foreground" scrub performed by Media Scrub Vdisk is typically faster than a background scrub performed by Vdisk Scrub. You can use a vdisk 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 vdisk

1. In the Configuration View panel, right-click a vdisk and select **Tools > Media Scrub Vdisk**.
2. Click **Start Media Scrub Utility**. A message confirms that the scrub has started.
3. Click **OK**. The panel shows the scrub's progress.

To abort a vdisk scrub

1. In the Configuration View panel, right-click a vdisk and select **Tools > Media Scrub Vdisk**.

NOTE: If the vdisk is being scrubbed but the Abort Media Scrub Utility button is grayed out, a background scrub is in progress. To stop the background scrub, disable the Vdisk Scrub option as described in ["Configuring background scrub for vdisks"](#) (page 54).

2. Click **Abort Media Scrub Utility**. A message confirms that the scrub has been aborted.
3. Click **OK**.

Removing a vdisk from quarantine

△ **CAUTION:** Carefully read this topic to determine whether to use the Dequarantine Vdisk panel to manually remove a vdisk from quarantine. The Dequarantine Vdisk panel should only be used as part of the emergency procedure to attempt to recover data and is normally followed by use of the CLI `trust` command. If a vdisk is manually dequarantined and does not have enough disks to continue operation, its status will change to OFFL and its data may or may not be recoverable through use of the `trust` command. It is recommended that you contact technical support for assistance in determining if the recovery procedure that makes use of the Dequarantine Vdisk panel and the `trust` command is applicable to your situation and for assistance in performing it. Also, see the help for the `trust` command.

To continue operation (that is, not go to quarantined status), a RAID-3 or RAID-5 vdisk can have only one inaccessible disk; a RAID-6 vdisk can have only one or two inaccessible disks; a RAID-10 or RAID-50 vdisk can have only one inaccessible disk per sub-vdisk. For example, a 16-disk RAID-10 vdisk can remain online (critical) with 8 inaccessible disks if one disk per mirror is inaccessible.

The system will automatically quarantine a vdisk having a fault-tolerant RAID level if one or more of its disks becomes inaccessible, or to prevent invalid ("stale") data that may exist in the controller from being written to the vdisk. Quarantine will not occur if a known-failed disk becomes inaccessible or if a disk becomes inaccessible after failover or recovery. The system will automatically quarantine an NRAID or RAID-0 vdisk to prevent invalid data from being written to the vdisk. If quarantine occurs because of an inaccessible disk, event 172 is logged. If quarantine occurs to prevent writing invalid data, event 485 is logged.

Examples of when quarantine can occur are:

- At system power-up, a vdisk has fewer disks online than at the previous power-up. This may happen because a disk is slow to spin up or because an enclosure is not powered up. The vdisk will be automatically dequarantined if the inaccessible disks come online and the vdisk status becomes FTOL (fault tolerant and online), or if after 60 seconds the vdisk status is QTCR or QTDN.
- During system operation, a vdisk loses redundancy plus one more disk; for example, three disks are inaccessible in a RAID-6 vdisk or two disks are inaccessible for other fault-tolerant RAID levels. The vdisk will be automatically dequarantined if after 60 seconds the vdisk status is FTOL, FTDN, or CRIT.

Quarantine isolates the vdisk from host access and prevents the system from changing the vdisk status to OFFL (offline). The number of inaccessible disks determines the quarantine status; from least to most severe:

- QTDN (quarantined with a down disk): The RAID-6 vdisk has one inaccessible disk. The vdisk is fault tolerant but degraded. If the inaccessible disks come online or if after 60 seconds from being quarantined the vdisk is QTCR or QTDN, the vdisk is automatically dequarantined.

- QTCR (quarantined critical): The vdisk is critical with at least one inaccessible disk. For example, two disks are inaccessible in a RAID-6 vdisk 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 vdisk is QTCR or QTDN, the vdisk is automatically dequarantined.
- QTOF (quarantined offline): The vdisk is offline with multiple inaccessible disks causing user data to be incomplete, or is an NRAID or RAID-0 vdisk.

When a vdisk is quarantined, its disks become write-locked, its volumes become inaccessible, and it is not available to hosts until it is dequarantined. If there are interdependencies between the quarantined vdisk's volumes and volumes in other vdisks, quarantine may temporarily impact operation of those other volumes. For example, if the quarantined vdisk contains the snap pool used for snapshot, volume-copy, or replication operations, quarantine may temporarily cause the associated master volume to go offline; a volume-copy or replication operation can also be disrupted if an associated volume (snap pool, source volume, or destination volume) goes offline. Depending on the operation, the length of the outage, and the settings associated with the operation, the operation may automatically resume when the vdisk is dequarantined or may require manual intervention. A vdisk can remain quarantined indefinitely without risk of data loss.

A vdisk is dequarantined when it is brought back online, which can occur in three ways:

- If the inaccessible disks come online, making the vdisk FTOL, the vdisk is automatically dequarantined.
- If after 60 seconds from being quarantined the vdisk is QTCR or QTDN, the vdisk is automatically dequarantined. The inaccessible disks are marked as failed and the vdisk status changes to CRIT (critical) or FTDN (fault tolerant with a down disk). If the inaccessible disks later come online, they are marked as LEFTOVR (leftover).
- The dequarantine command is used to manually dequarantine the vdisk. If the inaccessible disks later come online, they are marked as LEFTOVR (leftover). If event 485 was logged, use the dequarantine command only as specified by the event's recommended-action text to avoid data corruption or loss.

A quarantined vdisk can be fully recovered if the inaccessible disks are restored. Make sure that all disks are properly seated, that no disks have been inadvertently removed, and that no cables have been unplugged. Sometimes not all disks in the vdisk power up. Check that all enclosures have restarted after a power failure. If these problems are found and then fixed, the vdisk recovers and no data is lost.

If the inaccessible disks cannot be restored (for example, they failed), and the vdisk's status is FTDN or CRIT, and compatible spares are available, reconstruction will automatically begin.

If a replacement disk (reconstruct target) is inaccessible at power up, the vdisk becomes quarantined; when the disk is found, the vdisk is dequarantined and reconstruction starts. If reconstruction was in process, it continues where it left off.

NOTE: The only tasks allowed for a quarantined vdisk are Dequarantine Vdisk and Delete Vdisk. If you delete a quarantined vdisk and its inaccessible disks later come online, the vdisk will reappear as quarantined or offline and you must delete it again (to clear those disks).

To remove a vdisk from quarantine

1. In the Configuration View panel, right-click a quarantined vdisk and select **Tools > Dequarantine Vdisk**.
2. Click **Dequarantine Vdisk**. Depending on the number of disks that remain active in the vdisk, its health might change to Degraded (RAID 6 only) and its status changes to FTOL, CRIT, or FTDN. For status descriptions, see "[Vdisk properties](#)" (page 100).

Expanding a snap pool

By default, snap pools are configured to automatically expand when they become 90% full.

However, if a snap pool's policy is *not* set to Auto Expand and the snap pool is running out of free space, you can manually expand the snap pool.

For expansion to succeed, the vdisk must have free space and sufficient resources. Because expansion does not require I/O to be stopped, the snap pool can continue to be used during expansion.

To expand a snap pool

1. In the Configuration View panel, right-click a volume and select **Tools > Expand Snap Pool**.
2. In the main panel, specify the amount of free space to add to the snap pool.
3. Click **Expand Snap Pool**. If the task succeeds, the snap pool's size is updated in the Configuration View panel.

Checking links to a remote system

After a remote system has been added, you can check the connectivity between host ports in the local system and the remote system. A host port in the local system can only link to other host ports with the same host interface, such as Fibre Channel (FC), in a remote system. When you check links, this panel will show this information for each linked host port in the local system:

- The link type
- The ID of the port in the local system
- The ID of each accessible port in the remote system

If a host port is not shown then either:

- It is not linked
- Its link type is not supported by both systems

To check links to a remote system

1. In the Configuration View panel, right-click a remote system and select **Tools > Check Remote System Link**.
2. Click **Check Links**.

Checking local system links

You can check the connectivity between host ports in both controllers in the local system. A host port can only link to other ports with the same host interface. When you check links, this panel will show this information for each linked host port in both controllers:

- The link type
- The port ID
- The ID of each linked port in the local system

To check links in the local system

1. In the Configuration View panel, right-click the local system and select **Tools > Check Local System Link**.
2. Click **Check Links**.

Resetting or saving historical disk-performance statistics

Resetting historical disk-performance statistics

You can reset (clear) all historical performance statistics for all disks. When you reset historical statistics, an event will be logged and new data samples will continue to be stored every quarter hour.

To reset historical disk performance statistics

1. In the Configuration View panel, right-click the local system and select **Tools > Reset or Save Disk Performance Statistics**.
2. In the main panel, under the Reset Disk Performance Statistics heading, click **Reset**. A confirmation dialog appears.
3. Click **Yes** to continue; otherwise, click **No**. If you clicked Yes, a processing dialog appears. When processing is complete a success dialog appears.
4. Click **OK**.

Saving historical disk-performance statistics

You can download historical disk-performance statistics for all disks in the storage system. This task 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", ...
"2012-01-18 01:00:00", "disk_1.1", "PLV2W1XE", "2467917", ...
"2012-01-18 01:15:00", "disk_1.1", "PLV2W1XE", "2360042", ...
...
```

To save historical disk-performance statistics

1. In the Configuration View panel, right-click the local system and select **Tools > Reset or Save Disk Performance Statistics**.
2. In the main panel, under the Save Disk Performance Statistics heading, specify start and end dates and times to define the range of performance data to retrieve.
3. Click **Save**.





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 Reset or Save Disk Performance Statistics panel and retry the save operation.

4. When prompted to open or save the file, click **Save**.
 - If you are using Firefox and have a download directory set, the file `Disk_Performance.csv` is saved there.
 - 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.

5 Viewing system status

Viewing information about the system

In the Configuration View panel, right-click the system and select **View > Overview**. The System Overview table shows:

- Health.
 -  OK
 -  Degraded
 -  Fault
 -  Unknown
- Component. System, Enclosures, Disks, Vdisks, Volumes, Schedules, Configuration Limits, Versions, Snap Pools, Snapshots, Licensed Features.
- Count.
- Capacity.
- Storage Space. For descriptions of storage-space color codes, see "[About storage-space color codes](#)" (page 31).





Select a component to see more information about it.

NOTE: If the system is not working properly and you cannot determine why, you can restore its default configuration settings. You then can reconfigure the settings that are necessary to use the system. To restore defaults, in the CLI use the `restore defaults` command, as described in the CLI Reference Guide.

System properties

When you select System in the System Overview table, two tables display information about the system.

The System Information table shows:



- Health.
 -  OK
 -  Degraded
 -  Fault
 -  Unknown
- Health Reason. If the system's health is not OK, its Health Reason specifies that a subcomponent is unhealthy. In the System Overview table, notice which other components are unhealthy and view their properties as described in the following sections.
- System Name. User-defined system name.
- System Contact. User-defined system contact.
- System Location. User-defined system location.
- System Information. User-defined description of the system.
- Vendor Name.
- Product ID.
- Product Brand.
- SCSI Vendor ID.
- SCSI Product ID.
- Supported Locales. Languages supported by the system.

The System Redundancy table shows:

- Controller Redundancy Mode.
- Controller Redundancy Status.
- Controller A Status.
- Controller B Status.

Enclosure properties

When you select Enclosures in the System Overview table, a table displays the following information for each enclosure:




- Health.
 -  OK
 -  Degraded
 -  Fault
 -  Unknown

If an enclosure's health is not OK, select it in the Configuration View panel to view details about it.

- Enclosure ID.
- Enclosure WWN.
- Vendor.
- Model.
- Disk Slots. The quantity of disk slots in the enclosure.

Disk properties

When you select Disks in the System Overview table, a table shows:

- Health.
 -  OK
 -  Degraded
 -  Fault
 -  Unknown

If a disk's health is not OK, select it in the Configuration View panel to view details about it.

- Enclosure ID.
- Slot. The number of the slot the disk resides in.
- Serial Number.
- Vendor.
- Model.
- Revision.
- Type.
 - SAS: Enterprise SAS.
 - SAS MDL: Midline SAS.
 - sSAS: SAS SSD.

- How Used

Two values are listed together: the first is How Used and the second is Current Job. For example, for a disk used in a vdisk (VDISK) that is being scrubbed (VRSC), VDISKVRSC displays.





- How Used
 - AVAIL: Available.
 - FAILED: The disk is unusable and must be replaced. Reasons for this status include: excessive media errors; SMART error; disk hardware failure; unsupported disk.
 - GLOBAL SP: Global spare.

- LEFTOVR: Leftover.
- VDISK: Used in a vdisk.
- VDISK SP: Spare assigned to a vdisk.
- Current Job
 - DRSC: Disks in the vdisk are being scrubbed.
 - EXPD: The vdisk is being expanded.
 - INIT: The vdisk is being initialized.
 - RCON: The vdisk is being reconstructed.
 - VRFY: The vdisk is being verified.
 - VRSC: The vdisk is being scrubbed.
- Status.
 - 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.
 - Error: The disk is present but is 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.
- Size. Total size of the disk.

If there are no disks in the system, the table displays no data.

Vdisk properties

When you select Vdisks in the System Overview table, a table displays the following information for each vdisk:

- Health.
 -  OK
 -  Degraded
 -  Fault
 -  Unknown

If a vdisk's health is not OK, select it in the Configuration View panel to view details about it.
- Name. Vdisk name.
- Size. Total size of the vdisk.
- Free. Amount of free space remaining on the vdisk.
- RAID. RAID level.
- Status.
 - CRIT: Critical. The vdisk is online but isn't fault tolerant because some of its disks are down.
 - FTDN: Fault tolerant with down disks. The vdisk is online and fault tolerant, but some of its disks are down.
 - FTOL: Fault tolerant and online.
 - OFFL: Offline. Either the vdisk is using offline initialization, or its disks are down and data may be lost.
 - QTCR: Quarantined critical. The vdisk is offline and quarantined because at least one disk is missing. For instance, one disk is missing from a mirror or RAID 5.
 - QTDN: Quarantined with down disk. The RAID-6 vdisk has one inaccessible disk. The vdisk is fault tolerant but degraded.
 - QTOF: Quarantined offline. The vdisk is offline and quarantined because multiple disks are missing and user data is incomplete.
 - STOP: The vdisk is stopped.
 - UNKN: Unknown.
 - UP: Up. The vdisk is online and does not have fault-tolerant attributes

- Disk Type.
 - SAS: Enterprise SAS.
 - SAS MDL: Midline SAS.
 - sSAS: SAS SSD.

NOTE: In the Configuration View panel, if a vdisk contains more than one type of disk, its RAID-level label includes the suffix `-MIXED`.

Volume properties

When you select Volumes in the System Overview table, a table displays the following information for each volume:

- Name.
- Serial Number.
- Size. Total size of the volume.
- Vdisk Name. The name of the vdisk the volume resides on.

If no volumes exist, the table displays no data.

Snap-pool properties

When you select Snap Pools in the System Overview table, a table shows each snap pool's name, serial number, size, free space, master volumes, snapshots, and vdisk name.

If no snap pools exist, the table displays no data.

Snapshot properties

When you select Snapshots in the System Overview table, a table shows each snapshot's name; serial number; source volume; snap-pool name; amounts of snap data, unique data, and shared data; and vdisk name.

- Snap data is 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).
- Unique data is the amount of data that has been written to the snapshot since the last snapshot was taken. If the snapshot has not been written or is deleted, this value is zero bytes.
- Shared data is the amount of data that is potentially shared with other snapshots and the associated amount of space that will be freed if the snapshot is deleted. This represents the amount of data written directly to the snapshot. It also includes data copied from the source volume to the storage area for the oldest snapshot, since that snapshot does not share data with any other snapshot. For a snapshot that is not the oldest, if the modified data is deleted or if it had never been written to, this value is zero bytes.

If no snapshots exist, the table displays no data.

Schedule properties

When you select Schedules in the System Overview table, a table displays the following information for each schedule:

- Schedule Name.
- Schedule Specification. The start day and time of the schedule.
- Status. Expired or Active.
- Next Time. The next time the task is scheduled to run.

When you select a schedule, two tables display: the Schedule Details table and the Task Details table.

The Schedule Details table displays specifics about the schedule:

- Schedule Name.
- Schedule Specification. The start day and time of the schedule.
- Status. Expired or Active.

- Next Time. The next time the task is scheduled to run.

The Task Details table displays specifics about the task:

- Task Name.
- Task Type. Type of task assigned to run.
- Status. Outcome of the task.
- Task State. Specific information about task type.

When you select a task of type TakeSnapshot, a third table displays. The Retained Set table shows the name and serial number of each snapshot that the task has created and that is being retained.

If no schedules exist, the table displays no data.

To modify or delete scheduled tasks to suspend (disable) and resume (re-enable) DSD, use the Advanced Settings Disk panel; see "[Scheduling drive spin down for all disks](#)" (page 51).

Configuration limits

When you select Configuration Limits in the System Overview table, a table shows the Maximum Vdisks, Maximum Volumes, Maximum LUNs, Maximum Disks, and Number of Host Ports that the system supports. For a summary of the physical and logical limits of the storage system, see the system configuration limits topic in RAIDar help.

Licensed features

When you select Licensed Features in the System Overview table, a table shows the status of licensed features.

Version properties





When you select Versions in the System Overview table, a table shows the versions of firmware and hardware in each controller module.

- Storage Controller CPU Type.
- Bundle Version.
- Build Date.
- Storage Controller Code Version.
- Storage Controller Code Baselevel.
- Memory Controller FPGA Code Version.
- Storage Controller Loader Code Version.
- CAPI Version.
- Management Controller Code Version.
- Management Controller Loader Code Version.
- Expander Controller Code Version.
- CPLD Code Version.
- Hardware Version.
- Host Interface Module Version.
- Host Interface Module Model.
- Backplane Type.
- Host Interface Hardware (Chip) Version.
- Disk Interface Hardware (Chip) Version.
- SC Boot Memory Reference Code.

Viewing the system event log

In the Configuration View panel, right-click the system and select **View > Event Log**. The System Events panel shows the 100 most recent events that have been logged by either controller. All events are logged, regardless of event-notification settings. Click the buttons above the table to view all events, or only critical, warning, or informational events.

The event log table shows the following information:

- Severity.
 -  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.
- Time. Date and time when the event occurred, shown as *year-month-day hour:minutes:seconds*. Time stamps have one-second granularity.
- Event ID. An identifier for the event. The prefix A or B identifies the controller that logged the event.
- Code. An event code that helps you and support personnel diagnose problems. For event-code descriptions and recommended actions, see the Event Descriptions Reference Guide.
- Message. Brief information about the event. Click the message to show or hide additional information and recommended actions.





NOTE: If you are having a problem with the system or a vdisk, check the event log before calling technical support. Event messages might enable you to resolve the problem.

When reviewing events, do the following:

1. For any critical, error, or warning events, click the message to view additional information and recommended actions. This information also appears in the Event Descriptions Reference Guide.
Identify the primary events and any that might be the cause of the primary event. For example, an over-temperature event could cause a disk failure.
2. View the event log and locate other critical/error/warning events in the sequence for the controller that reported the event.
Repeat this step for the other controller if necessary.
3. Review the events that occurred before and after the primary event.
During this review you are looking for any events that might indicate the cause of the critical/error/warning event. You are also looking for events that resulted from the critical/error/warning event, known as secondary events.
4. Review the events following the primary and secondary events.
You are looking for any actions that might have already been taken to resolve the problems reported by the events.

Viewing information about all vdisks

In the Configuration View panel, right-click **Vdisks** and select **View > Overview**. The Vdisks Overview table shows:





- Health.
 -  OK
 -  Degraded
 -  Fault
 -  Unknown
- Component.
- Count. Number of components.
- Capacity. Total capacity of the component.
- Storage Space. Amount of space on the component. For descriptions of storage-space color codes, see ["About storage-space color codes" \(page 31\)](#).

The Vdisks table shows more information about each vdisk.

- Health.
- Name. Vdisk name.
- Size. Total storage space in the vdisk.
- Free. Available space in the vdisk.
- RAID. RAID level of the vdisk and all of its volumes.
- Status.
 - CRIT: Critical. The vdisk is online but isn't fault tolerant because some of its disks are down.
 - FTDN: Fault tolerant with a down disk. The vdisk is online and fault tolerant, but some of its disks are down.
 - FTOL: Fault tolerant and online.
 - OFFL: Offline. Either the vdisk is using offline initialization, or its disks are down and data may be lost.
 - QTCR: Quarantined critical. The vdisk is critical with at least one inaccessible disk. For example, two disks are inaccessible in a RAID-6 vdisk 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 vdisk is QTCR or QTDN, the vdisk is automatically dequarantined.
 - QTDN: Quarantined with a down disk. The RAID-6 vdisk has one inaccessible disk. The vdisk is fault tolerant but degraded. If the inaccessible disks come online or if after 60 seconds from being quarantined the vdisk is QTCR or QTDN, the vdisk is automatically dequarantined.
 - QTOF: Quarantined offline. The vdisk is offline with multiple inaccessible disks causing user data to be incomplete, or is an NRAID or RAID-0 vdisk.
 - STOP: The vdisk is stopped.
 - UNKN: Unknown.
 - UP: Up. The vdisk is online and does not have fault-tolerant attributes.
- Disk Type.
 - SAS: Enterprise SAS.
 - SAS MDL: Midline SAS.
 - sSAS: SAS SSD.
- Preferred Owner. Controller that owns the vdisk and its volumes during normal operation.
- Current Owner. Either the preferred owner during normal operation or the partner controller when the preferred owner is offline.
- Disks. Quantity of disks in the vdisk.
- Spares. Quantity of dedicated spares in the vdisk.

Viewing information about a vdisk

In the Configuration View panel, right-click a vdisk and select **View > Overview**. The Vdisks Overview table shows:





- Health.
 -  OK
 -  Degraded
 -  Fault
 -  Unknown
- Component. Vdisk, disks, volumes.
- Count.
- Capacity. For descriptions of storage-space color codes, see ["About storage-space color codes" \(page 31\)](#).

Select a component to see more information about it. When the Vdisk component is selected, you can view properties or historical performance statistics.

NOTE: Failure of a disk in the vdisk causes the Vdisk and Disks components to have Degraded health. Because tables displayed when the Disks component is selected exclude failed disks, those tables will show fewer disks than the Disk component's Count value.

Vdisk properties

When you select Vdisk in the Vdisk Overview table and click the **Properties** tab, the Properties for Vdisk table shows:

- Health.
 -  OK
 -  Degraded
 -  Fault
 -  Unknown
- Health Reason. If Health is not OK, this field shows the reason for the health state.
- Health Recommendation. If Health is not OK, this field shows recommended actions to take to resolve the health issue.
- Name. Vdisk name.
- Size. Total storage space in the vdisk.
- Free. Available space in the vdisk.
- Current Owner. Either the preferred owner during normal operation or the partner controller when the preferred owner is offline.
- Preferred Owner. Controller that owns the vdisk and its volumes during normal operation.
- Serial Number. Vdisk serial number.
- RAID. RAID level of the vdisk and all of its volumes.
- Disks. Quantity of disks in the vdisk.
- Spares. Quantity of dedicated spares in the vdisk.
- Chunk Size.
 - For RAID levels except NRAID, RAID 1, and RAID 50, the configured chunk size for the vdisk.
 - For NRAID and RAID 1, chunk size has no meaning and is therefore shown as not applicable (N/A).
 - For RAID 50, the vdisk chunk size calculated as: $configured_chunk_size \times (subvdisk_members - 1)$. For a vdisk configured to use 32-KB chunk size and 4-disk sub-vdisks, the value would be 96K (32KB x 3).
- Created. Date and time when the vdisk was created.
- Minimum Disk Size. Capacity of the smallest disk in the vdisk.

- Status.
 - CRIT: Critical. The vdisk is online but isn't fault tolerant because some of its disks are down.
 - FTDN: Fault tolerant with a down disk. The vdisk is online and fault tolerant, but some of its disks are down.
 - FTOL: Fault tolerant and online.
 - OFFL: Offline. Either the vdisk is using offline initialization, or its disks are down and data may be lost.
 - QTCR: Quarantined critical. The vdisk is critical with at least one inaccessible disk. For example, two disks are inaccessible in a RAID-6 vdisk 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 vdisk is QTCR or QTDN, the vdisk is automatically dequarantined.
 - QTDN: Quarantined with a down disk. The RAID-6 vdisk has one inaccessible disk. The vdisk is fault tolerant but degraded. If the inaccessible disks come online or if after 60 seconds from being quarantined the vdisk is QTCR or QTDN, the vdisk is automatically dequarantined.
 - QTOF: Quarantined offline. The vdisk is offline with multiple inaccessible disks causing user data to be incomplete, or is an NRAID or RAID-0 vdisk.
 - STOP: The vdisk is stopped.
 - UNKN: Unknown.
 - UP: Up. The vdisk is online and does not have fault-tolerant attributes.
- Current Job. If a utility is running on the vdisk, this field shows the utility's name and progress.
 - Disk Scrub: Disks in the vdisk are being scrubbed.
 - Expand: The vdisk is being expanded.
 - Initialize: The vdisk is being initialized.
 - Reconstruct: The vdisk is being reconstructed.
 - Verify: The vdisk is being verified.
 - Media Scrub: The vdisk is being scrubbed.
- Drive Spin Down Vdisk Enable. Shows whether drive spin down is enabled or disabled for this vdisk.

A second table displays information about unhealthy components. If all components are healthy, this table displays the text, "There is no data for your selection."

NOTE: In the Configuration View panel, if a vdisk contains more than one type of disk, its RAID-level label includes the suffix `-MIXED`.

Vdisk performance

When you select Vdisk in the Vdisk Overview table and click the **Performance** tab, the Performance Statistics panel shows three graphs of historical performance statistics for the vdisk: Data Transferred, Data Throughput, and Average Response Time. Data samples are taken every quarter hour and the graphs represent up to 50 samples.

To specify a time range of samples to display, set the start and end values and click **Update**. The system determines whether the number of samples in the time range exceeds the number of samples that can be displayed (50), requiring aggregation. To determine this, the system divides the number of samples in the specified time range by 50, giving a quotient and a remainder. If the quotient is 1, the 50 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 50 so all 4 samples are displayed.
- Example 2: A 15-hour range includes 60 samples. 60 divided by 50 gives a quotient of 1 and a remainder of 10. Therefore, the newest 50 samples will be displayed and the oldest 10 samples will be excluded.
- Example 3: A 30-hour range includes 120 samples. 120 divided by 50 gives a quotient of 2 and a remainder of 20. Therefore, each two newest samples will be aggregated into one sample for display and the oldest 20 samples will be excluded.

If aggregation is required, the system aggregates samples for each disk in the vdisk (as described in "[Disk performance](#)" (page 113)) and then aggregates the resulting values as follows:

- For a count statistic such as data transferred, the aggregated values are added to produce the value of the aggregated sample.
- For a rate statistic such as data throughput, the aggregated values are added and then are divided by their combined interval (seconds per sample multiplied by the number of samples).

The system will change the time settings to match the times of the oldest and newest samples displayed. The graphs are updated each time you click either the Performance tab or the Update button.

- For the vdisk, the Data Transferred graph shows the amounts of data read and written and the combined total over the sampling time period. The base unit is bytes.
- For the vdisk, the Data Throughput graph shows the rates at which data are read and written and the combined total over the sampling time period. The base unit is bytes per second.
- For each disk in the vdisk, the Average Response Time graph shows the average response times for reads and writes over the sampling time period. The base unit is microseconds. To view the graph's legend, which identifies the color-coding for each disk, select **Show Legend**.

 **TIP:** If you specify a time range, it is recommended to specify a range of 12 hours or less.

To view performance data for an individual disk, use the Enclosure Overview panel ([page 110](#)). To view live (non-historical) performance statistics for one more vdisks, in the CLI use the `show vdisk-statistics` command.

NOTE: Values for the amount of data transferred and for data throughput appear to be much higher in historical output than in live output. This is caused by a difference in the way that historical and live values are calculated.

Live values are calculated based on the vdisk as viewed from the controller cache perspective. In the live statistics, performance numbers are obtained by accounting for when data is written from cache to disk or is read from disk to cache.

Historical data is obtained by using the summation of the disk statistics for the disks in the vdisk. The historical vdisk data shows transfers to and from the disks in the vdisk that include the overhead of any RAID transfers as well as any host activity.

Because I/Os from the RAID engine are included, values for the historical data appear higher than the numbers for the live data.

Disk properties

When you select Disks in the Vdisk Overview table, a Disk Sets table and enclosure view appear. The Disk Sets table shows:

- Total Space. Total storage space in the vdisk, followed by a color-coded measure of how the space is used.
- Type. For RAID 10 or RAID 50, the sub-vdisk that the disk is in; for other RAID levels, the disk's RAID level; or SPARE.
- Disk Type.
 - SAS: Enterprise SAS.
 - SAS MDL: Midline SAS.
 - sSAS: SAS SSD.
- Disks. Quantity of disks in the vdisk or sub-vdisk.
- Size. Total capacity of the disks in the vdisk or sub-vdisk.

The enclosure view table has two tabs. The Tabular tab shows:

- Health.
 -  OK
 -  Degraded
 -  Fault
 -  N/A (if the disk is spun down)
 -  Unknown

If the disk's health is not OK, view health details in the Enclosure Overview panel ([page 110](#)).
- Name. System-defined disk name using the format *Disk-enclosure-number.disk-slot-number*.
- Type.
 - SAS: Enterprise SAS.
 - SAS MDL: Midline SAS.
 - sSAS: SAS SSD.
- State. Shows how the disk is used:
 - If the disk is in a vdisk, its RAID level
 - AVAIL: Available
 - FAILED: The disk is unusable and must be replaced. Reasons for this status include: excessive media errors; SMART error; disk hardware failure; unsupported disk.
 - SPARE: Spare assigned to a vdisk
 - GLOBAL SP: Global spare
 - LEFTOVR: Leftover

This also shows any job running on the disk:

 - DRSC: The disk is being scrubbed
 - EXPD: The vdisk is being expanded
 - INIT: The vdisk is being initialized
 - RCON: The vdisk is being reconstructed
 - VRFY: The vdisk is being verified
 - VRSC: The vdisk is being scrubbed
- Size. Disk capacity.
- Enclosure. Name of the enclosure containing the disk.
- Serial Number. Disk serial number.
- Status. Up (operational) or Not Present.

The Graphical tab shows the locations of the vdisk's disks in system enclosures and each disk's Health and State.

Volume properties

When you select Volumes in the Vdisk Overview table, the Volumes table shows:

- Name. Volume name.
- Serial Number. Volume serial number.
- Size. Volume size.
- Vdisk Name. The name of the vdisk containing the volume.

Snap-pool properties

When you select Snap Pools in the Vdisk Overview table, the Snap Pools table shows:

- The snap pool's name, serial number, size, and free space.
- The quantity of master volumes and snapshots associated with the snap pool.
- The name of the vdisk containing the snap pool.

Viewing information about a volume

In the Configuration View panel, right-click a volume and select **View > Overview**. The Volume Overview table shows:

- Component. Volume or Maps.
- Count. The quantity of mappings for the volume.
- Capacity. The capacity of the volume.
- Storage Space. The space usage of the volume. For descriptions of storage-space color codes, see ["About storage-space color codes" \(page 31\)](#).
- Replication Addresses. The quantity of replication addresses for the volume.
- Replication Images. The quantity of replication images for the volume.

Select a component to see more information about it.

Volume properties

When you select Volume in the Volume Overview table, the Properties for *Volume* table shows:

- Vdisk Name. Name of the vdisk that the volume is in.
- Name. Volume name.
- Size. Volume size.
- Preferred Owner. Controller that owns the vdisk and its volumes during normal operation.
- Current Owner. Either the preferred owner during normal operation or the partner controller when the preferred owner is offline.
- Serial Number. Volume serial number.
- Cache Write Policy. Write-back or Write-through. See ["Using write-back or write-through caching" \(page 22\)](#).
- Cache Optimization. Standard or No-mirror. See ["Optimizing read-ahead caching" \(page 23\)](#).
- Read Ahead Size. See ["Optimizing read-ahead caching" \(page 23\)](#).
- Type. Standard volume, master volume, or snapshot.
- Progress. If the volume is being created by a volume-copy operation, the percent complete.
- Snapshot. Specifies whether the selected volume is a snapshot.
- Health. OK, Degraded, Fault, or Unknown.
Health Reason. If Health is not OK, this field shows the reason for the health state.
- Health Recommendation. If Health is not OK, this field shows the recommended actions to take to resolve the health issue.

For a *local* primary or secondary volume, the Replication Properties for Volume table shows:

- Name. Replication volume name.
- Serial Number. Replication volume serial number.
- Status. Replication volume status:
 - Initializing: The initial (full) replication to the volume is in progress.
 - Online: The volume is online and is consistent with the last replicated image.
 - Inconsistent: The volume is online but is in an inconsistent state. A full replication is required to initialize it.
 - Replicating: The volume is online and replication is in progress.
 - Replicate-delay: The volume is online but the in-progress replication has been temporarily delayed; a retry is occurring.
 - Suspended: The volume is online but the in-progress replication has been suspended.
 - Offline: The volume cannot be accessed or is unusable due to an error.
 - Establishing proxy: The volume is establishing a proxy connection to a remote volume. This will occur when a detached secondary volume is reattached and is re-establishing a connection with the primary system in preparation for replication.
 - Detached: The volume is detached for removal.
- Status-Reason. More information about the status value, or N/A for Online status.

- Monitor. Replication volume monitoring status:
 - OK: Communication to the remote volume is successfully occurring on the FC or iSCSI network.
 - Failed: Communication to the remote volume has failed because of an FC or iSCSI network issue or because the remote volume has gone offline.
- Location. Local or Remote.
- Primary Volume Name. Primary volume name. If the replication set has a primary-volume conflict, all associated primary volumes are displayed.
- Primary Volume Serial Number. Primary volume serial number. If the replication set has a primary-volume conflict, all associated primary volumes are displayed.
- Primary Volume Status. Primary volume status: Online, Offline, Conflict, or N/A.
- Maximum Number of Queued Images. Number of replication images to consider when determining the next image to replicate. Used only if the On Collision parameter is set to Oldest.
- Maximum Retry Time (Seconds). Amount of time in seconds that the replication volume should retry a replication operation on any specific image when errors occur. Used only if the On Error parameter is set to Retry.
- On Error. Error policy to invoke when errors occur during the replication process: Retry or Suspend.
- Link Type. Type of ports used to link the primary and secondary volumes: FC or iSCSI.
- On Collision. Collision policy used to determine the next image to replicate when multiple replication images are queued: Newest or Oldest.
- Monitor Interval. Interval in seconds at which the primary volume should query the secondary volume.
- Priority. Priority of the replication process on the replication volume: Low, Medium, or High.
- Connection Status.
 - Not Attempted. Communication has not been attempted to the remote volume.
 - Online. The volumes in the replication set have a valid connection but communication is not currently active.
 - Active. Communication is currently active to the remote volume.
 - Offline. No connection is available to the remote system.
- Connection Time. Date and time of the last communication with the remote volume, or N/A.

Mapping properties

When you select Maps in the Volume Overview table, the Maps for *Volume* table shows:

- Type. Explicit or Default. Settings for an explicit mapping override the default mapping.
- Host ID. WWPN or IQN.
- Host Name. User-defined nickname for the host.
- Ports. Controller host ports through which the volume is mapped to the host.
- LUN. Volume identifier presented to the host.
- Access. Volume access type: read-write, read-only, no-access (masked), or not-mapped.

Schedule properties

If any schedules exist for this volume, when you select the Schedules component, the Schedules table shows each schedule's name, specification, status, next run time, task type, task status, and task state. For the selected schedule, two tables appear.

The Schedule Details table shows:

- Schedule Name. Schedule name.
- Schedule Specification. The schedule's start time and recurrence or constraint settings.
- Status.
 - Uninitialized: Schedule is not yet ready to run.
 - Ready: Schedule is ready to run.
 - Suspended: Schedule is suspended.

- Expired: Schedule has expired.
- Invalid: Schedule is invalid.
- Next Time.

The Task Details table shows different properties depending on the task type. Properties shown for all task types are:

- Task Name. Task name.
- Task Type. ReplicateVolume, ResetSnapshot, TakeSnapshot, or VolumeCopy.
- Status.
 - Uninitialized: Task is not yet ready to run.
 - Ready: Task is ready to run.
 - Active: Task is running.
 - Error: Task has an error.
 - Invalid: Task is invalid.
- Task State. Current step of task processing. Steps vary by task type.
- Source Volume. Name of the volume to snap, copy, or replicate.
- Source Volume Serial. Source volume serial number.
- Destination Vdisk. Name of the destination vdisk for a volume copy.
- Destination Vdisk Serial. Destination vdisk serial number.
- Prefix. Label that identifies snapshots, volume copies, or replication images created by this task.
- Count. Number of snapshots to retain with this prefix. When a new snapshot exceeds this limit, the oldest snapshot with the same prefix is deleted.
- Last Created. Name of the last snapshot, volume copy, or replication image created by the task.
- Last Used Snapshot. For a task whose replication mode is last-snapshot, the name of the last snapshot used for replication.
- Snapshot Name. Name of the snapshot to reset.
- Snapshot Serial. Snapshot serial number.
- Mode. Replication mode:
 - new-snapshot: Replicate a new snapshot of the primary volume.
 - last-snapshot: Replicate the most recent existing snapshot of the primary volume.

For a TakeSnapshot task, the Retained Set table shows the name and serial number of each snapshot that the task has taken and is retaining.

Viewing information about a snapshot

In the Configuration View panel, right-click a snapshot and select **View > Overview**. The Snapshot Overview table shows:

- The capacity and space usage of the snapshot
- The quantity of mappings for the snapshot
- The quantity of task schedules for the snapshot

For descriptions of storage-space color codes, see ["About storage-space color codes" \(page 31\)](#).

Select a component to see more information about it.

Snapshot properties

When you select the Snapshot component, the Properties for Snapshot table shows:

- Vdisk Name.
- Serial Number. Snapshot serial number.
- Name. Snapshot name.
- Creation Date/Time.
- Status.

- Status-Reason.
- Source Volume. Name of the volume that the snapshot was taken of.
- Snap-pool Name.
- Snap Data. 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).
- UniqueData. The amount of data that has been written to the snapshot since the last snapshot was taken. If the snapshot has not been written or is deleted, this value is zero bytes.
- SharedData. The amount of data that is potentially shared with other snapshots and the associated amount of space that will be freed if the snapshot is deleted. This represents the amount of data written directly to the snapshot. It also includes data copied from the source volume to the storage area for the oldest snapshot, since that snapshot does not share data with any other snapshot. For a snapshot that is not the oldest, if the modified data is deleted or if it had never been written to, this value is zero bytes.
- Priority. Retention priority for the snapshot, based on the snapshot attributes and the user-defined priority for the snapshot type.
- User Priority. User-defined priority for the snapshot type.
- Type.
 - Standard snapshot: Snapshot of a master volume that consumes a snapshot license.
 - Standard snapshot (DRM): A temporary standard snapshot created from a replication snapshot for the purpose of doing a test failover for disaster recovery management (DRM).
 - Replication snapshot: For a primary or secondary volume, a snapshot that was created by a replication operation but is not a sync point.
 - Replication snapshot (Replicating): For a primary volume, a snapshot that is being replicated to a secondary system.
 - Replication snapshot (Current sync point): For a primary or secondary volume, the latest snapshot that is copy-complete on any secondary system in the replication set.
 - Replication snapshot (Common sync point): For a primary or secondary volume, the latest snapshot that is copy-complete on all secondary systems in the replication set.
 - Replication snapshot (Old Common sync point): For a primary or secondary volume, a common sync point that has been superseded by a new common sync point.
 - Replication snapshot (Only sync point): For a primary or secondary volume, the only snapshot that is copy-complete on any secondary system in the replication set.
 - Replication snapshot (Queued): For a primary volume, a snapshot associated with a replication operation that is waiting for a previous replication operation to complete.
 - Replication snapshot (Awaiting replicate): For a primary volume, a snapshot that is waiting to be replicated to a secondary system.

Mapping properties

When you select the Maps component, the Maps for *Volume* table shows:

- Type. Explicit or Default. Settings for an explicit mapping override the default mapping.
- Host ID. WWPN or IQN.
- Name. Host name.
- Ports. Controller host ports through which the volume is mapped to the host.
- LUN. Volume identifier presented to the host.
- Access. Volume access type: read-write, read-only, no-access (masked), or not-mapped.

Schedule properties

If any schedules exist for the snapshot, when you select the Schedules component, the Schedules table shows information about each schedule. For the selected schedule, the Schedule Details table shows:

- Schedule Name.
- Schedule Specification.
- Schedule Status.
- Next Time.
- Task Type.
- Task Status.
- Task State.
- Source Volume.
- Source Volume Serial.
- Prefix.
- Count.
- Last Created.

Viewing information about a snap pool

In the Configuration View panel, right-click a snap pool and select **View > Overview**. The Snap Pool Overview table shows:

- The capacity and space usage of the snap pool
- The quantity of volumes using the snap pool
- The quantity of snapshots in the snap pool

For descriptions of storage-space color codes, see ["About storage-space color codes" \(page 31\)](#).

NOTE: The process of freeing space associated with deleted snapshots occurs more slowly when the system is operating write-through cache mode than in write-back cache mode. Therefore, there will be a delay between deleting the snapshots and when their used space is shown as available space in the Snap Pool Overview panel.

Select a component to see more information about it.

Snap-pool properties

When you select the Snap Pool component, two tables appear. The first table shows the snap pool's name, serial number, size (total capacity), vdisk name, and free space, the number of snapshots in the snap pool, and its status. The status values are:

- Available: The snap pool is available for use.
- Offline: The snap pool is not available for use, as in the case where its disks are not present.
- Corrupt: The snap pool's data integrity has been compromised; the snap pool can no longer be used.

The second table shows the snap pool's threshold values and associated policies. Three thresholds are defined:

- Warning: The snap pool is moderately full. When this threshold is reached, an event is generated to alert the administrator. The default value is 75%.
- Error: The snap pool is nearly full and unless corrective action is taken, snapshot data loss is probable. When this threshold is reached, an event is generated to alert the administrator and the associated snap-pool policy is triggered. The default is 90%.
- Critical: The snap pool is 98% full and data loss is imminent. When this threshold is reached, an event is generated to alert the administrator and the associated snap-pool policy is triggered.

The following policies are defined:

- **Auto Expand:** Automatically expand the snap pool by the indicated expansion-size value. This is the default policy for the Error threshold.
If the snap pool's space usage reaches the percentage specified by its error threshold, the system will log Warning event 230 and will try to automatically expand the snap pool by the snap pool's expansion-size value. If the snap pool cannot be expanded because there is not enough available space in its vdisk, the system will log Warning event 444 and will automatically delete the oldest snapshot that is not a current sync point.
- **Delete Oldest Snapshot:** Delete the oldest snapshot.
- **Delete Snapshots:** Delete all snapshots. This is the default policy for the Critical threshold.
- **Halt Writes:** Halt writes to all master volumes and snapshots associated with the snap pool.
- **Notify Only:** Generates an event to notify the administrator. This is the only policy for the Warning threshold.
- **No Change:** Take no action.

NOTE: The policies Delete Oldest Snapshot and Delete Snapshots do not apply business logic to the delete decision and may delete snapshots that are mounted/presented/mapped or modified. You may set retention priorities for a snap pool as a way of suggesting that some snapshots are more important than others, but these priorities do not ensure any specific snapshot is protected.

For details about setting snap-pool thresholds and policies, see the CLI Reference Guide.

Volume properties

When you select the Client Volumes component, a table shows the name, serial number, size, vdisk name, and vdisk serial number for each volume using the snap pool.

Snapshot properties

When you select the Resident Snapshots component, a table shows each snapshot's name; serial number; amounts of snap data, unique data, and shared data; and status (Available or Unavailable).

Snap data is 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).

Unique data is the amount of data that has been written to the snapshot since the last snapshot was taken. If the snapshot has not been written or is deleted, this value is zero bytes.

Shared data is the amount of data that is potentially shared with other snapshots and the associated amount of space that will be freed if the snapshot is deleted. This represents the amount of data written directly to the snapshot. It also includes data copied from the source volume to the storage area for the oldest snapshot, since that snapshot does not share data with any other snapshot. For a snapshot that is not the oldest, if the modified data is deleted or if it had never been written to, this value is zero bytes.

Viewing information about all hosts

In the Configuration View panel, right-click **Hosts** and select **View > Overview**. The Hosts table shows the quantity of hosts configured in the system.

For each host, the Hosts Overview table shows the following details:

- **Host ID.** WWPN or IQN.
- **Name.** User-defined nickname for the host.
- **Discovered.** If the host was discovered and its entry was automatically created, Yes. If the host entry was manually created, No.
- **Mapped.** If volumes are mapped to the host, Yes; otherwise, No.
- **Profile.**
 - **Standard:** Default profile.

- Host Type.
 - If the host was discovered and its entry was automatically created, its host-interface type: FC; iSCSI; SAS.
 - If the host entry was manually created: Undefined.

Viewing information about a host

In the Configuration View panel, right-click a host and select **View > Overview**. The Host Overview table shows:

- Host properties
- The quantity of mappings for the host

Select a component to see more information about it.

Host properties

When you select Host in the Host Overview table, the Properties for *Host* table shows:

- Host ID. WWPN or IQN.
- Name. User-defined nickname for the host.
- Discovered. If the host was discovered and its entry was automatically created, Yes. If the host entry was manually created, No.
- Mapped. If volumes are mapped to the host, Yes; otherwise, No.
- Profile.
 - Standard: Default profile.
- Host Type.
 - If the host was discovered and its entry was automatically created, its host-interface type: FC; iSCSI; SAS.
 - If the host entry was manually created: Undefined.

Mapping properties

When you select Maps in the Host Overview table, the Maps for *Host* table shows:

- Type. Explicit or Default. Settings for an explicit mapping override the default mapping.
- Name. Volume name.
- Serial Number. Volume serial number.
- Ports. Controller host ports through which the volume is mapped to the host.
- LUN. Volume identifier presented to the host.
- Access. Volume access type: read-write, read-only, no-access (masked), or not-mapped.

Viewing information about an enclosure

In the Configuration View panel, right-click an enclosure and select **View > Overview**. You can view information about the enclosure and its components in a front or rear graphical view, or in a front or rear tabular view.





- Front Graphical. Shows a graphical view of the front of each enclosure and its disks.
- Front Tabular. Shows a tabular view of each enclosure and its disks.
- Rear Graphical. Shows a graphical view of components at the rear of the enclosure.
- Rear Tabular. Shows a tabular view of components at the rear of the enclosure.

Tabular views are initially sorted by the Name property.

In any of these views, select a component to see more information about it. Components vary by enclosure model. If any components are unhealthy, a table at the bottom of the panel identifies them. When a disk is selected, you can view properties or historical performance statistics.






Enclosure properties

When you select an enclosure, a table shows:

- Health.
 -  OK
 -  Degraded
 -  Fault
 -  Unknown
- Health Reason. If Health is not OK, this field shows the reason for the health state.
- Health Recommendation. If Health is not OK, this field shows recommended actions to take to resolve the health issue.
- Status.
- Enclosure ID.
- Vendor.
- Model.
- Disk Slots.
- Enclosure WWN.
- Midplane Serial Number.
- Part Number.
- Manufacturing Date.
- Manufacturing Location.
- Revision.
- EMP A Revision. Firmware revision of the Enclosure Management Processor in controller module A's Expander Controller.
- EMP B Revision. Firmware revision of the Enclosure Management Processor in controller module B's Expander Controller.
- EMP A Bus ID.
- EMP B Bus ID.
- EMP A Target ID.
- EMP B Target ID.
- Midplane Type.
- Enclosure Power (watts).
- PCIe 2-Capable. Shows whether the enclosure is capable of using PCI Express version 2.

Disk properties

When you select a disk and click the **Properties** tab, a table shows:




- Health.
 -  OK
 -  Degraded
 -  Fault
 -  N/A (if the disk is spun down)
 -  Unknown
- Health Reason. If Health is not OK, this field shows the reason for the health state.
- Health Recommendation. If Health is not OK, this field shows recommended actions to take to resolve the health issue.

- Status.
 - 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.
 - Error: The disk is present but is 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.

For an SSD, the Front Tabular view also shows the percentage of disk life remaining from 100% to 0%. When the value decreases to 20%, event 502 is logged with Informational severity. Event 502 is logged again with Warning severity when the value decreases to 5%, 2%, and 0%.

- Enclosure ID.
- Slot.
- How Used.
 - AVAIL: Available.
 - FAILED: The disk is unusable and must be replaced. Reasons for this status include: excessive media errors; SMART error; disk hardware failure; unsupported disk.
 - GLOBAL SP: Global spare.
 - LEFTOVR: Leftover.
 - VDISK: Used in a vdisk.
 - VDISK SP: Spare assigned to a vdisk.
- Type.
 - SAS: Enterprise SAS.
 - SAS MDL: Midline SAS.
 - sSAS: SAS SSD.
- Vendor.
- Model.
- Size.
- Speed (kr/min).
- Transfer Rate. The data transfer rate in Gbit/s. Some 6-Gbit/s disks might not consistently support a 6-Gbit/s transfer rate. If this happens, the controller automatically adjusts transfers to those disks to 3 Gbit/s, increasing reliability and reducing error messages with little impact on system performance. This rate adjustment persists until the controller is restarted or power-cycled.
- Revision. Disk firmware revision number.
- Serial Number.
- Current Job.
 - DRSC: Disks in the vdisk are being scrubbed.
 - EXPD: The vdisk is being expanded.
 - INIT: The vdisk is being initialized.
 - RCON: The vdisk is being reconstructed.
 - VRFY: The vdisk is being verified.
 - VRSC: The vdisk is being scrubbed.
- SMART. Shows whether Self-Monitoring Analysis and Reporting Technology is enabled. For more information, see ["Configuring SMART" \(page 50\)](#).
- Current Owner. For the disk's vdisk, either the preferred owner during normal operation or the partner controller when the preferred owner is offline.
- Drive Spin Down Count. How many times the disk has been spun down.

- SSD Life Remaining.
 - For an SSD this property shows the percentage of disk life remaining, represented by a color-coded bar graph:

Life remaining	Bar color	Disk health
100–20%		OK
19–5%		OK
4–1%		Degraded
0%	(No bar)	Critical or Fault

When the value decreases to 20%, event 502 is logged with Informational severity. Event 502 is logged again with Warning severity when the value decreases to 5%, 2%, and 0%.

- For a non-SSD this property shows N/A.

Disk performance

When you select a disk and click the **Performance** tab, a table shows eight graphs of historical performance statistics for the disk. Data samples are taken every quarter hour and the graphs represent up to 50 samples. By default, the graphs show the newest 50 samples.

To specify a time range of samples to display, set the start and end values and click **Update**. The system determines whether the number of samples in the time range exceeds the number of samples that can be displayed (50), requiring aggregation. To determine this, the system divides the number of samples in the specified time range by 50, giving a quotient and a remainder. If the quotient is 1, the 50 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 50 so all 4 samples are displayed.
- Example 2: A 15-hour range includes 60 samples. 60 divided by 50 gives a quotient of 1 and a remainder of 10. Therefore, the newest 50 samples will be displayed and the oldest 10 samples will be excluded.
- Example 3: A 30-hour range includes 120 samples. 120 divided by 50 gives a quotient of 2 and a remainder of 20. Therefore, each two newest samples will be aggregated into one sample for display and the oldest 20 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/s.

- 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.

The system will change the time settings to match the times of the oldest and newest samples displayed. The graphs are updated each time you click either the Performance tab or the Update button.

- Data Transferred. Shows the amounts of data read and written and the combined total over the sampling time period. The base unit is bytes.
- Data Throughput. Shows the rates at which data are read and written and the combined total over the sampling time period. The base unit is bytes/s.
- I/O. Shows the numbers of reads and writes and the combined total over the sampling time period.
- IOPS. Shows numbers of reads and writes per second and the combined total over the sampling time period.
- Average Response Time. Shows the average response times for reads and writes and the combined average over the sampling time period. The base unit is microseconds.





- Average I/O Size. Shows the average sizes of reads and writes and the combined average over the sampling time period. The base unit is bytes.
- Disk Error Counters. Shows the number of disk errors over the sampling time period.
- Average Queue Depth. Shows the average number of pending I/O operations that are being serviced over the sampling time period. This value represents periods of activity only and excludes periods of inactivity.

 **TIP:** If you specify a time range, it is recommended to specify a range of 12 hours or less.

To view summary performance data for a vdisk, use the Vdisk Overview panel as described on [page 100](#). To view live (non-historical) performance statistics for one more disks, in the CLI use the `show disk-statistics` command.




Power supply properties

When you select a power supply, a table shows:

- Health.
 -  OK
 -  Degraded
 -  Fault
 -  Unknown
- Health Reason. If Health is not OK, this field shows the reason for the health state.
- Health Recommendation. If Health is not OK, this field shows recommended actions to take to resolve the health issue.
- Status.
- Model.
- Vendor.
- Location.
- Serial Number.
- Revision.
- Part Number.
- Manufacturing Date.
- Manufacturing Location.

Controller module properties



When you select a controller module, a table shows:

- Health.
 -  OK
 -  Fault
 -  Unknown
- Health Reason. If Health is not OK, this field shows the reason for the health state.
- Health Recommendation. If Health is not OK, this field shows recommended actions to take to resolve the health issue.
- Status.
- Controller ID.
- Description.
- CPLD Version.
- Storage Controller Code Version.
- Model.
- Storage Controller CPU Type.

- Serial Number.
- Part Number.
- Position.
- Hardware Version.
- Revision.
- Manufacturing Date.
- Manufacturing Location.





Controller module: network port properties

When you select a network port, a table shows:

- Health.
 -  OK
 -  Degraded
- Health Reason. If Health is not OK, this field shows the reason for the health state.
- MAC Address.
- Addressing Mode.
- IP Address.
- Gateway.
- Subnet Mask.

Controller module: FC host port properties




When you select an FC host port, a table shows:

- Health.
 -  OK
 -  Degraded
 -  Fault
 -  N/A
- Health Reason. If Health is not OK, this field shows the reason for the health state.
- Status.
 - Up: The port is cabled and has an I/O link.
 - Warning: Not all of the port's PHYs are up.
 - 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.
- Ports. The port ID, which is the controller ID and port number.
- Media.
 - FC(L): Fibre Channel-Arbitrated Loop (public or private).
 - FC(P): Fibre Channel Point-to-Point.
 - FC(-): Fibre Channel disconnected.
- Target ID. The port WWN.
- Configured Speed. Auto, 4Gb, 8Gb, or 16Gb (Gbit/s).
- Actual Speed. Actual link speed in Gbit/s, or blank if not applicable.
- Configured Topology.
 - PTP: Fibre Channel Point-to-Point.
 - Loop: Fibre Channel-Arbitrated Loop (public or private).
- Primary Loop ID. Primary loop ID, or blank if not applicable.

- SFP Status.
 - OK
 - Not present: No SFP is inserted in this port.
 - Not compatible: The SFP in this port is not qualified for use in this system. When this condition is detected, event 464 is logged.
 - Incorrect protocol: The SFP protocol does not match the port protocol. When this condition is detected, event 464 is logged.
- Part Number. The SFP part number.
- Supported Speeds. The link speeds that the SFP supports, in Gbit/s.

Controller module: iSCSI host port properties






When you select an iSCSI host port, a table shows:

- Health.
 -  OK
 -  Degraded
 -  Fault
 -  N/A
 -  Unknown
- Health Reason. If Health is not OK, this field shows the reason for the health state.
- Status.
 - Up: The port is cabled and has an I/O link.
 - Warning: Not all of the port's PHYs are up.
 - 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.
- Ports. The port ID, which is the controller ID and port number.
- Media. iSCSI.
- Target ID. The port IQN.
- Configured Speed. Auto; the link speed is auto-negotiated.
- Actual Speed. Actual link speed in Gbit/s, or blank if not applicable.
- IP Version. The IP version: IPv4 or IPv6.
- MAC. The port's MAC address.
- IP Address. For IPv4 or IPv6, assigned port IP address.
- Netmask. For IPv4, subnet mask for assigned port IP address.
- Gateway. For IPv4, gateway for assigned port IP address.
- Default Router. For IPv6, default router for assigned port IP address.
- Link-Local Address. For IPv6, the link-local address that is automatically generated from the MAC address and assigned to the port.
- SFP Status.
 - OK
 - Not present: No SFP is inserted in this port.
 - Not compatible: The SFP in this port is not qualified for use in this system. When this condition is detected, event 464 is logged.
 - Incorrect protocol: The SFP protocol does not match the port protocol. When this condition is detected, event 464 is logged.
- 10G Compliance. The SFP's 10G compliance code, if supported.
- Cable Length. The link length (in meters) that is supported by the SFP while operating in compliance with applicable standards for the cable type, or 0 if this information is not provided by the SFP manufacturer.

- Cable Technology. Shows whether the SFP supports active or passive cable technology, or N/A if this information is not provided by the SFP manufacturer.
- Ethernet Compliance. The SFP's Ethernet compliance code, if supported.
- Part Number. The SFP part number.






Controller module: SAS host port properties

When you select a SAS host port, a table shows:

- Health.
 -  OK
 -  Degraded
 -  Fault
 -  N/A
 -  Unknown
- Health Reason. If Health is not OK, this field shows the reason for the health state.
- Health Recommendation. If Health is not OK, this field shows recommended actions to take to resolve the health issue.
- Status.
 - Up: The port is cabled and has an I/O link.
 - Warning: Not all of the port's PHYs are up.
 - 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.
- Ports. The port ID, which is the controller ID and port number.
- Media. SAS.
- Target ID. The port WWN.
- Configured Speed. Blank; not applicable for SAS.
- Actual Speed. Auto; the link speed is auto-negotiated.



Controller module: expansion port properties

When you select an expansion (Out) port, a table shows:

- Health.
 -  OK
 -  Degraded
 -  Fault
 -  N/A
 -  Unknown
- Health Reason. If Health is not OK, this field shows the reason for the health state.
- Health Recommendation. If Health is not OK, this field shows recommended actions to take to resolve the health issue.
- Status.
- Name.

Controller module: CompactFlash properties





When you select a CompactFlash card in the Rear Tabular view, a table shows:

- Health.
 -  OK
 -  Fault
- Health Reason. If Health is not OK, this field shows the reason for the health state.

- Health Recommendation. If Health is not OK, this field shows recommended actions to take to resolve the health issue.
- Status.
- Cache Flush.
 - Enabled: If the controller loses power, it will automatically write cache data to the CompactFlash card. Cache flush is normally enabled, but is temporarily disabled during controller shut down.
 - Disabled: Cache flush is disabled.






Drive enclosure: I/O module properties

When you select an I/O module, a table shows:

- Health.
 -  OK
 -  Degraded
 -  Fault
 -  Unknown
- Health Reason. If Health is not OK, this field shows the reason for the health state.
- Status.
- Controller ID.






I/O module: In port properties

When you select an In port, a table shows:

- Health.
 -  OK
 -  Degraded
 -  Fault
 -  N/A
 -  Unknown
- Health Reason. If Health is not OK, this field shows the reason for the health state.
- Health Recommendation. If Health is not OK, this field shows recommended actions to take to resolve the health issue.
- Status.
- Name.

I/O module: Out port properties

When you select an Out port, a table shows:

- Health.
 -  OK
 -  Degraded
 -  Fault
 -  N/A
 -  Unknown
- Health Reason. If Health is not OK, this field shows the reason for the health state.
- Health Recommendation. If Health is not OK, this field shows recommended actions to take to resolve the health issue.
- Status.
- Name.

Viewing information about a remote system

In the Configuration View panel, right-click a remote system and select **View > Overview**. The System Information table shows:

- The username and network-port IP addresses that are configured on the local system to access the remote system. The configured password is not shown.
- Information such as the system name, location, and status that is read from the remote system.

To sign in to the remote system, click one of its IP address links.

6 Using AssuredRemote to replicate volumes

About the AssuredRemote replication feature

AssuredRemote™ replication is a licensed feature for disaster recovery. This feature performs asynchronous (batch) replication of block-level data from a volume on a local storage system to a volume that can be on the same system or on a second, independent system. This second system can be located at the same site as the first system or at a different site.

TIP: For guidelines about setting up and using the AssuredSnap, AssuredCopy, and AssuredRemote features singly or together, see the document *AssuredSAN 3000 Series Using Data Protection Software*.

A typical replication configuration involves these physical and logical components:

- A host connected to a local storage system, which is networked via FC or iSCSI ports to a remote storage system as described in installation documentation.
- *Remote system.* A management object on the local system that enables the MCs in the local system and in the remote system to communicate and exchange data.
- *Replication set.* Associated master volumes that are enabled for replication and that typically reside in two physically or geographically separate storage systems. These volumes are also called replication volumes.
- *Primary volume.* The volume that is the source of data in a replication set and that can be mapped to hosts. For disaster recovery purposes, if the primary volume goes offline, a secondary volume can be designated as the primary volume. The primary volume exists in a primary vdisk in the primary system.
- *Secondary volume.* The volume that is the destination for data in a replication set and that is not accessible to hosts. For disaster recovery purposes, if the primary volume goes offline, a secondary volume can be designated as the primary volume. The secondary volume exists in a secondary vdisk in a secondary system.
- *Replication snapshot.* A special type of snapshot that preserves the state of data of a replication set's primary volume as it existed when the snapshot was created. For a primary volume, the replication process creates a replication snapshot on both the primary system and, when the replication of primary-volume data to the secondary volume is complete, on the secondary system. Replication snapshots are unmappable and are not counted toward a license limit, although they are counted toward the system's maximum number of volumes. A replication snapshot can be exported to a regular, licensed snapshot.
- *Replication image.* A conceptual term for replication snapshots that have the same image ID in the primary and secondary systems. These synchronized snapshots contain identical data and can be used for disaster recovery.

Replication process overview

As a simplified overview of the remote-replication process, it can be configured to provide a single point-in-time replication of volume data or a periodic delta-update replication of volume data.

The periodic-update process has multiple steps. At each step, matching snapshots are created: in the primary system, a replication snapshot is created of the primary volume's current data; this snapshot is then used to copy new (delta) data from the primary volume to the secondary volume; then in the secondary system, a matching snapshot is created for the updated secondary volume. This pair of matching snapshots establishes a replication sync point and these sync points are used to continue the replication process.

The following figure illustrates three replication sets in use by two hosts:

- The host in New York is mapped to and updates the Finance volume. This volume is replicated to the system in Munich.
- The host in Munich is mapped to and updates the Sales and Engineering volumes. The Sales volume is replicated from System 2 to System 3 in the Munich data center. The Engineering volume is replicated from System 3 in Munich to System 1 in New York.




Figure 4 Intersite and intrasite replication sets

Remote replication uses snapshot functionality to track the data to be replicated and to determine the differences in data updated on the master volume, minimizing the amount of data to be transferred.

In order to perform a replication, a snapshot of the primary volume is taken, creating a point-in-time image of the data. This point-in-time image is then replicated to the secondary volume by copying the data represented by the snapshot using a transport medium such as TCP/IP (iSCSI) or Fibre Channel. The first replication copies all data from the primary volume to the secondary volume; subsequent replications use sparse snapshots.

Replication snapshots are retained for both the primary volume and the secondary volume. When a matching pair of snapshots is retained for both volumes, the matching snapshots are referred to as *replication sync points*. The two snapshots (one on each volume) are used together as a synchronization reference point, minimizing the amount of data to transfer. The two snapshots in a replication sync point are assigned the same *image ID*, which uniquely identifies that the data in those snapshots are from the same point-in-time image and are block-for-block identical.

When a replication snapshot is created from a standard snapshot, while that snapshot remains present the replication snapshot's total data represented is zero bytes. This behavior occurs because the snapshot data remains associated with the standard snapshot and there is no data specifically associated with the replication snapshot. If the standard snapshot is deleted, its data becomes associated with (is preserved by) the replication snapshot and the replication snapshot's size changes to reflect the size of the deleted snapshot.

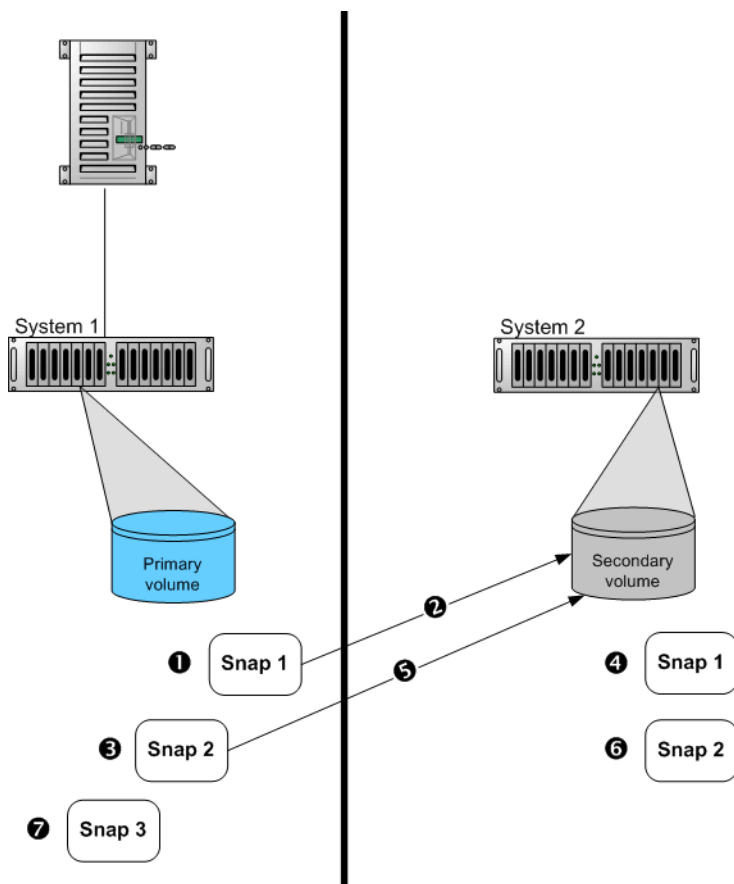
An added benefit of using snapshots for replication is that these snapshots can be kept and restored later in the event of a non-hardware failure, such as virus attack. Since the replication source is a snapshot, any writes performed on the primary volume after the snapshot is taken are not replicated by that task. This gives you more control over what is contained in each replication image.

NOTE: Because replication is not synchronous (continuous), data in a secondary volume is only as current as the last replication that completed successfully. Replications can be performed manually or scheduled.

NOTE: Snapshot operations are I/O-intensive. Every write to a unique location in a master volume after a snapshot is taken will cause an internal read and write operation to occur in order to preserve the snapshot data. If you intend to create snapshots of, create volume copies of, or replicate volumes in a vdisk, ensure that the vdisk contains no more than four master volumes, snap pools, or both. For example: 2 master volumes and 2 snap pools; 3 master volumes and 1 snap pool; 4 master volumes and 0 snap pools.

Replication actions

The following figure illustrates actions that occur during a series of replications from System 1 to System 2.



- 1 Take initial snapshot and initiate replication.
- 2 Initial replication consists of a full data copy.
- 3 Take second snapshot and request replication. This can be taken while the initial replication is in progress.
- 4 Snapshot taken on secondary volume. This is the first replication sync point.
- 5 When the initial replication is complete, the second replication automatically starts. Only the data changed since Snap 1 is replicated.
- 6 Second snapshot taken on secondary volume. This is the second sync point.
- 7 Other snapshots can be taken and replication initiated on the primary volume while replication is in progress. These snapshots are queued waiting for prior replications to complete. These replication snapshots will not become sync points until their replications are complete.

Figure 5 Actions that occur during a series of replications


The figure above illustrates initial, delta, and queued replications:

- Initial replication: When the first replication is initiated, a snapshot of the primary volume is taken and every block of data is then copied to the secondary volume. When the copy is complete, the first snapshot is taken on the secondary volume, creating the first sync point. This sync point can be used to determine the delta data from that sync point to a later snapshot. Actions 1–4 are the initial replication.
- Delta replications: Delta data is the “list” of 64-KB blocks that differs between the last snapshot replicated and the next snapshot to be replicated. This delta data is then replicated from the replication snapshot on the primary

volume to the secondary volume. Once the initial replication has completed, all future replications for that replication set will be delta replications so long as sync points are maintained. Action 5 is a delta replication.

- Queued replications: New replications can be initiated while other replication snapshots are in the process of being replicated. This enables you to take snapshots at specific intervals while other replications are ongoing. Note that a replication that is initiated while another to the same secondary volume is ongoing will be queued, and will not begin to transfer data until the prior one completes. In action 3, Snap 2 is queued while Snap 1 is being replicated; in action 7, Snap 3 is queued while Snap 2 is being replicated.

An in-progress replication can be suspended, either manually by a user or automatically if a network error occurs. If you want the replication to continue, you must manually resume it; or, if you want to cancel the replication, you can abort it.

 **IMPORTANT:** For a replication to begin, the controller that owns the secondary volume must have a link to the controller that owns the primary volume. This link must be of the type specified by the link-type parameter supplied during replication set creation or modification. If all links to the controller that owns the primary volume fail, but links remain between its partner controller and the controller that owns the secondary volume, replications currently in progress or queued may continue, but their progress may not be reported correctly. If the controller that owns the secondary volume loses all links to both controllers of the primary system, then the replications will suspend and progress will be updated appropriately; links from the partner controller of the controller that owns the secondary volumes are not considered for use. Replications that enter the suspended state must be resumed manually.

Performing initial replication locally or remotely

When you set up replication for a volume, you specify to use a secondary volume in a vdisk in either the local (primary) system or a remote (secondary) system.

- If the speed of the initial replication is most important, specify a vdisk that is owned by the same controller as the primary volume's vdisk in the local system.
After replication is set up, you can perform the initial replication and then physically move the vdisk containing the secondary volume and its snap pool into a remote system. Moving a vdisk involves using RAIDar to detach the secondary volume and stop its vdisk, removing the vdisk's disks or enclosure, transporting the disks or enclosure to the remote location, inserting the disks or enclosure into the remote system, and using RAIDar to restart the vdisk and reattach the secondary volume. If the secondary volume's snap pool is in a different vdisk then that vdisk must also be stopped, moved, and restarted.
- If ease of setup is most important, specify a vdisk owned by either controller in a remote system. After replication is set up, you can start replication.

In either case, you must specify the link type to be used for replication between the primary and secondary systems and you cannot change this setting for the life of the replication set.

You can only select a vdisk that has enough available space for replication. For details, see the following topic.

Criteria for selecting a vdisk to contain a secondary volume

When setting up replication for a volume that will become the primary volume in a replication set, you have the option to select an existing vdisk in which to create the secondary volume.

The `vdisk-selection` option only lists vdisks that have sufficient available space for replication, and that do not contain a volume with a conflicting name (`xprimary-volume-name`) or a snap pool with a conflicting name (`sprimary-volume-name`). The system calculates the required space for the secondary volume (the reserve size) and its snap pool as follows:

- The snap-pool size will be either 20% of the primary volume's size or 6 GB, whichever is larger.
- The reserve size is calculated as follows:
 - If the primary volume and the snap pool are each less than 500 GB, the reserve will be the same size as the primary volume.
 - If the primary volume is larger than 500 GB, the reserve size will be the maximum, 500 GB.
 - If the snap pool is larger than 500 GB, the reserve will be the same size as the snap pool.

- The required space in the vdisk is calculated as follows:
 - If the combined size of the primary volume and the reserve is less than the combined size of the primary volume and the snap pool, the required space is the combined size of the primary volume and the snap pool.
 - If the combined size of the primary volume and the reserve is larger than the combined size of the primary volume and the snap pool, the required space is the combined size of the primary volume and the reserve.

The following table shows examples of how much available space a vdisk must have in order to be shown by the vdisk option. If you want to replicate a volume whose size is not shown, you can use the above calculations to determine how much available space the secondary vdisk must have.

Table 12 Available space required for a vdisk to be selectable to contain a secondary volume

Primary volume size (GB)	Available space required in vdisk (GB)	Primary volume size (GB)	Available space required in vdisk (GB)	Primary volume size (GB)	Available space required in vdisk (GB)
100	200	1100	1600	2100	2600
200	400	1200	1700	2200	2700
300	600	1300	1800	2300	2800
400	800	1400	1900	2400	2900
500	1000	1500	2000	2500	3000
600	1100	1600	2100	2600	3120
700	1200	1700	2200	2700	3240
800	1300	1800	2300	2800	3360
900	1400	1900	2400	2900	3480
1000	1500	2000	2500	3000	3600

Remote replication disaster recovery

Replication can continue in the event of system faults such as:

- Temporary communication failure. Remote replication will retry replication operations according to user-configured policies.
- Controller failure. In a dual-controller system, failover will occur and the surviving controller will take over replication processing until controller recovery occurs.
- Disk or power supply failure.

If a disaster causes the primary volume to become inaccessible, you can set the secondary volume to be the primary volume so that volume can be mapped to hosts. Disaster recovery requires user intervention because decisions must be made based on the data content of replication volumes and their snapshots.

1. Synchronize the secondary volume to a replication snapshot, preferably a replication sync point. Any data written to the primary volume since the last-completed replication will not be available.
2. After synchronization, set the secondary volume to be the new primary volume.
3. Map the new primary volume to hosts, as was the original primary volume.




Figure 6 Example of primary-volume failure

If the original primary volume becomes accessible, you can set it to be the primary volume again as described in the following process overview:

1. Take a snapshot of the original primary volume. This preserves the volume's current data state for later comparison with the new primary volume.
2. Remove the volume's mappings.
3. Set the original primary volume to be a secondary volume.
4. Replicate any data written to the new primary volume to the original primary volume (now a secondary volume). This can be done as one or more replications. On the final replication, halt host access to the primary volume to ensure that all data has been transferred to the secondary volume.
5. Set the secondary volume (the original primary volume) to be the new primary volume.
6. You can now mount/present/map the snapshot taken in [step 1](#) and compare it with the new primary volume to identify any data discrepancies and try to recover any data from the snapshot that would otherwise be lost. For example, you could use host file-system tools to find any files modified since a certain time, or for a database you could export any differing records from the snapshot and re-enter them into the current database.

For details, see the procedure to change the primary volume back to the original primary volume in ["Changing the primary volume for a replication set"](#) (page 135).

Remote replication licensing

The AssuredRemote and AssuredSnap features are separately licensed. AssuredRemote can operate without AssuredSnap being enabled; however, to get the most out of AssuredRemote, it is recommended to enable both features. Normally, replication snapshots are not accessible to hosts. However, if AssuredSnap is enabled, a replication snapshot can be exported for use as a standard snapshot and will count toward the snapshot license limit.


Related topics

- ["Installing a license"](#) (page 41)
- [Adding](#) (page 55) or [deleting](#) (page 55) a remote system
- ["Checking links to a remote system"](#) (page 90)
- ["Using the Replication Setup Wizard"](#) (page 127)
- [Replicating a volume](#) (page 129) or a snapshot (page 130)
- [Detaching](#) (page 132) and [reattaching](#) (page 134) a secondary volume
- [Stopping](#) (page 133) and [restarting](#) (page 133) a vdisk


- Suspending ([page 131](#)), resuming ([page 132](#)), or aborting ([page 132](#)) a replication
- "Exporting a replication image to a snapshot" ([page 135](#))
- "Changing the primary volume for a replication set" ([page 135](#))
- "Viewing replication properties, addresses, and images for a volume" ([page 136](#))
- "Viewing information about a replication image" ([page 138](#))

Using the Replication Setup Wizard

If the system is licensed to use remote replication, you can use the Replication Setup Wizard to prepare to replicate an existing volume to another vdisk in the local system or to a remote system. Before using this wizard, read the documentation for your product to learn about replication. Then plan the storage systems, replication mode, and volumes you want to use for the replication.

The wizard guides you through the following steps. For each step you can view help by clicking the help icon  in the wizard panel. As you complete steps they are highlighted at the bottom of the panel. If you cancel the wizard at any point, no changes are made.

- Select the primary volume, which is an existing volume or snapshot to replicate.
- Specify whether the replication mode will be local or remote. If the replication will be to a remote system that has not already been added to the local system, you can add it. To do so, you must know the user name and password of a user with the Manage role on that system and the system's IP address.
- Specify the secondary volume. You can select an existing replication-prepared volume or specify to create a volume in an existing vdisk that has sufficient available space for the replicated data.
- Confirm changes and apply them.

 **IMPORTANT:** Before starting this procedure, if you intend to use CHAP to authenticate iSCSI login requests between the local system and a remote system, do the following:

- Create a one-way CHAP record on each system. On the local system, the CHAP record must refer to the node name of the remote system. On the remote system, the CHAP record must refer to the node name of the local system. Both records must use the same secret. (Mutual CHAP is not used between storage systems. CHAP records' mutual fields can be set but are not used.) To create a CHAP record, see "[Configuring CHAP](#)" ([page 76](#)).
- After the CHAP records are created, enable CHAP on the primary system, the secondary system, or both. To enable CHAP, see "[Changing host interface settings](#)" ([page 47](#)).

If both records don't exist or don't use the same secret, replication-set creation will fail.

Step 1: Starting the wizard

1. In the Configuration View panel, right-click the system and select **Wizards > Replication Setup Wizard**. The wizard panel appears.
2. Click **Next** to continue.

Step 2: Selecting the primary volume

Select the volume whose data you want to replicate. If the volume has at least one snapshot, you can select a snapshot to be the replication source.

To select the primary volume

1. Set the options:
 - Select the vdisk that contains the volume to replicate. Only vdisks that contain at least one volume are listed for selection.
 - Select the volume to replicate. Only volumes that are not already part of a replication set are listed for selection.
2. Click **Next** to continue.

Step 3: Selecting the replication mode

Select the replication mode, which specifies whether the secondary volume is in the local system or a remote system. If you want to replicate to a remote system that hasn't already been added to the local system, you can add it.

To replicate within the local system

1. Select **Local Replication**.
2. Although it is recommended to check host-port links between controllers in the local system, if you already know the status of links you can clear the **Check Links** checkbox to skip this task.
3. Click **Next** to continue. If Check Links is selected and there are no links between the controllers, a message appears, and only vdisks and volumes that are owned by the same controller as the primary volume will appear in the next step.

To replicate to a remote system

1. Select **Remote Replication**.
2. In the Remote System list, look for the remote system that you want to use.
 - If you find the system, select it and continue with [step 5](#).
 - If you don't find it, add it as described in [step 3](#).
3. To add a remote system, in the Add new Remote System area:
 - a. Enter the IP address of a network port on the remote system.
 - b. Enter the user name of a user with a Manage role on the remote system.
 - c. Enter that user's password.
 - d. Click **Add Remote System**. If task succeeds, the new remote system appears in the Remote System list and is selected.
4. Although it is recommended to check host-port links between the two systems, this can take up to 3 minutes, so if you already know the status of links you can clear the **Check Links** checkbox to skip this task.
5. Click **Next** to continue. If Check Links is selected and there are no links to the remote system, a message appears and you cannot proceed. For a CNC system, if only one link type is up, only that link type will appear in the next step.

Step 4: Selecting the secondary volume

Specify the secondary volume. You can either select an existing vdisk in which to create the secondary volume, or select an existing replication-prepared volume to be the secondary volume.

To specify the secondary volume

1. Either:
 - Select **Create new volume on vdisk** and select an existing vdisk in which to create the secondary volume. For an explanation of the criteria that determines which vdisks are listed for selection, see "[Criteria for selecting a vdisk to contain a secondary volume](#)" (page 124).
 - Select **Use existing replication-prepared volume** and select an existing replication-prepared volume to be the secondary volume. Only replication-prepared volumes that are exactly the same size in blocks as the primary volume are listed for selection.
2. Select the link type used between the two systems.
3. Click **Next** to continue.

Step 5: Confirming replication settings

Confirm that the values 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 setting changes and finish the wizard.


Replicating a volume

If the system is licensed to use remote replication, you can create a replication set that uses the selected volume as the primary volume, and to immediately start or schedule replication. The primary volume can be a standard volume or a master volume.

To create a replication set you must select a secondary system and a secondary vdisk or volume. The secondary system can be the local system, or a remote system added by using the Add Remote System panel. When using RAIDar it is recommended to select a secondary vdisk and let the secondary volume be created automatically, instead of selecting an existing secondary volume. For a secondary (replication-prepared) volume to be available for selection, it must be exactly the same size in blocks as the primary volume, and that is difficult to ensure, especially with maximum-size volumes.

You can select the local system if you intend to create the replication set on the local system and then physically move the secondary vdisk's disks (or enclosure) to a remote system. Otherwise, select a remote system that you've added to the local system.

 **TIP:** A best practice is to schedule no more than three volumes to start replicating at the same time, and for those replications to recur no less than 60 minutes apart. If you schedule more replications to start at the same time, or schedule replications to start more frequently, some scheduled replications may not have time to complete.

 **IMPORTANT:** Before starting this procedure, if you intend to use CHAP to authenticate iSCSI login requests between the local system and a remote system, do the following:

- Create a one-way CHAP record on each system. On the local system, the CHAP record must refer to the node name of the remote system. On the remote system, the CHAP record must refer to the node name of the local system. Both records must use the same secret. (Mutual CHAP is not used between storage systems. CHAP records' mutual fields can be set but are not used.) To create a CHAP record, see ["Configuring CHAP" \(page 76\)](#).
- After the CHAP records are created, enable CHAP on the primary system, the secondary system, or both. To enable CHAP, see ["Changing host interface settings" \(page 47\)](#).

If both records don't exist or don't use the same secret, replication-set creation will fail.

NOTE: If replication requests are sent to a secondary system whose temporary replication license has expired, the requests are queued but are not processed, and the secondary system reports event 472. If this condition occurs, check for this event in the event log, event-notification emails, and SNMP traps. To continue using replication, purchase a permanent replication license.

To create a replication set and optionally start or schedule replication

1. In the Configuration View panel, right-click a volume and select **Provisioning > Replicate Volume**.
2. In the main panel, set the destination options:
 - Secondary System. Select a storage system to replicate the volume to.
 - Secondary Volume. Select either an existing vdisk in which to create the secondary volume, or an existing replication-prepared volume to be the secondary volume. For an explanation of the criteria that determines which vdisks are listed for selection, see ["Criteria for selecting a vdisk to contain a secondary volume" \(page 124\)](#).
3. Select the link type used between the two systems.
4. If you want to start replication now:
 - a. Select the **Initiate Replication** and **Now** options.
 - b. Optionally change the default replication image name. A name is case sensitive; cannot already exist in a vdisk; cannot include a comma, double quote, or backslash; and can have a maximum of 20 bytes.
 - c. Continue with [step 7](#).

5. If you want to schedule replication:
 - a. Select the **Initiate Replication** and **Scheduled** options.
 - b. Set the options:
 - Replication image prefix. Optionally change the default prefix to identify images created by this schedule. The prefix is case sensitive; cannot include a comma, double quote, or backslash; and can have a maximum of 14 bytes.
 - Replication Mode. Specifies whether to replicate a new snapshot of the volume to the remote system, or to replicate the last (most recent existing) snapshot of the volume to the remote system.
 - Replication images to Retain. Select the number of replication images to retain for both the primary volume and the secondary volume. When the task runs, the retention count is compared with the number of existing replication images:
 - Whether the retention count has been reached or not, a new replication image is created.
 - If the retention count has been reached, the volume's oldest replication image that was created by this schedule and is neither being replicated, nor a current sync point, nor a queued snapshot, is deleted.
 - If there is more than one queued snapshot, only the oldest queued snapshot is retained. It is retained to serve as the source for the next scheduled replication to create a replication image from.
 - Start Schedule. 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. Date must use the format *yyyy-mm-dd*. 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.
 - Recurrence. Specify the interval at which the schedule should run. The default and minimum interval is 30 minutes.
 - Time Constraint. Specify a time range within which the schedule should run.
 - Date Constraint. Specify days when the schedule should run. Ensure that this constraint includes the Start Schedule date.
 - End Schedule. Specify when the schedule should stop running.
 - c. Continue with [step 7](#).
6. If you do not want to start or schedule replication, clear the **Initiate Replication** checkbox. The replication set will still be created and you can replicate the volume at a later time.
7. Click **Apply**. Within a couple of minutes, the replication set is created and the following changes occur in the Configuration View panel:
 - Under the primary vdisk:
 - The selected primary volume changes to a master volume, and is designated as a Primary Volume.
 - If the secondary volume is on a remote system, the secondary volume appears under the primary volume.
 - If a replication was performed, under both the primary volume and the secondary volume a replication image appears.
 - If not already present, the primary volume's snap pool appears.
 - Under the secondary vdisk:
 - The secondary volume appears.
 - If the primary volume is on a remote system, the primary volume appears under the secondary volume.
 - If a replication was performed, under both the primary volume and the secondary volume a replication image appears.
 - If not already present, the secondary volume's snap pool appears.

Replicating a snapshot

If the system is licensed to use remote replication, you can replicate an existing, primary snapshot that is mapped to a host. You can only replicate a snapshot of a volume that is already part of a replication set.

If the selected snapshot hasn't already been replicated to a secondary volume, each replication volume in the replication set is requested to replicate the snapshot data. Only snapshot preserved data is replicated; snapshot modified data is not replicated.

To replicate a snapshot

1. In the Configuration View panel, right-click a snapshot and select **Provisioning > Replicate Snapshot**.
2. In the main panel, optionally change the default replication image name. A name is case sensitive; cannot already exist in a vdisk; cannot include a comma, double quote, or backslash; and can have a maximum of 20 bytes.
3. Click **Initiate Replication**. In a few seconds, the following changes occur in the Configuration View panel:
 - Under both the primary volume and the secondary volume, a replication image appears.

Removing replication from a volume

If the system is licensed to use remote replication and you no longer want to replicate a volume, you can remove its replication set. When a replication set is removed:

- A rollback is automatically performed to the latest available snapshot on the secondary volume to ensure that data is consistent.
- Replication volumes associated with the replication set are converted to master volumes.
- Any replication images associated with the replication volumes are converted to standard snapshots. Snapshots are converted regardless of the number of snapshots allowed by the system's license.
- There is no longer a relationship between the volumes or their snapshots in the two vdisks.

To remove replication from a volume

1. In the Configuration View panel, right-click a local primary volume and select **Provisioning > Remove Replication Set**.
2. In the main panel, click **Remove Replication Set**. A confirmation dialog appears.
3. Click **Remove** to continue; otherwise, click **Cancel**. If you clicked Remove, a processing dialog appears. A message indicates whether the task succeeded or failed.
4. Click **OK**. If the task succeeded, the following changes occur in the Configuration View panel:
 - Under the primary vdisk:
 - The primary volume's designation is changed from Primary Volume to Volume
 - The secondary volume is removed
 - Any replication images are replaced by snapshots
 - Under the secondary vdisk:
 - The secondary volume's designation is changed from Secondary Volume to Volume
 - The primary volume is removed
 - Any replication images are replaced by snapshots

NOTE: Normally, if you want to remove a replication set you must select its primary volume. However, if the primary volume is inaccessible, you can set the secondary volume to be the primary volume (as described in ["Changing the primary volume for a replication set"](#) (page 135)) and then perform a Remove Replication Set operation.

Suspending replication

If the system is licensed to use remote replication, you can suspend the current replication operation for a selected a replication volume. You must perform this task on the system that owns the secondary volume. Once suspended, the replication must be resumed or aborted to allow the replication volume to resume normal operation.

To suspend replication

1. In the Configuration View panel, right-click a local replication volume and select **Provisioning > Suspend Replication**.
2. In the main panel, click **Suspend Replication**. A message indicates whether the task succeeded or failed.

Resuming replication

If the system is licensed to use remote replication, you can resume a suspended replication operation for a selected replication volume. You must perform this task on the system that owns the secondary volume.

To resume replication

1. In the Configuration View panel, right-click a local replication volume and select **Resume Replication**.
2. In the main panel, click **Resume Replication**. A message indicates whether the task succeeded or failed.

Aborting replication

If the system is licensed to use remote replication, you can abort the current replication operation for the selected replication volume. The current replication may be running or suspended. You must perform this task on the system that owns the secondary volume.

To abort replication

1. In the Configuration View panel, right-click a local replication volume and select **Provisioning > Abort Replication**.
2. In the main panel, click **Abort Replication**. A message indicates whether the task succeeded or failed.

Detaching a secondary volume

When using the replication feature, if you chose to create a replication set's primary and secondary volumes in the primary system, you can perform the initial replication and then physically move the secondary volume's vdisk into the secondary system.

The process to move a secondary volume is:

1. In the system where the secondary volume resides:
 - a. Detach the secondary volume.
 - b. If the secondary volume's vdisk contains other secondary volumes, detach those volumes.
 - c. Stop the secondary volume's vdisk. For details see ["Stopping a vdisk" \(page 133\)](#).
 - d. If the secondary volumes' snap pools are in other vdisks, stop those vdisks.
 - e. Move the vdisks into the secondary system. This system must support the link type that the replication set is configured to use. For example, if the replication set's link type is configured to use FC links, the secondary system must have FC ports.
2. In the secondary system:
 - a. Start the snap pools' vdisks. For details see ["Starting a vdisk" \(page 133\)](#).
 - b. Start the secondary volumes' vdisks.
 - c. Reattach the secondary volumes. For details see ["Reattaching a secondary volume" \(page 134\)](#).

Detached volumes remain associated with their replication sets but are not updated with replication data or with replication control information.

NOTE:

- It is recommended that the vdisk that you are moving contains only secondary volumes and their snap pools. You are allowed to move other volumes along with secondary volumes and their snap pools, but be sure that you are doing so intentionally.
 - If you intend to move a vdisk's enclosure and you want to allow I/O to continue to the other enclosures, it is best if it is at the end of the chain of connected enclosures. If the enclosure is in the middle of the chain, the enclosures must be cabled with no single point of failure, so that removing the enclosure does not prevent communication between other enclosures.
-

To detach a secondary volume

1. In the Configuration View panel, right-click the secondary volume and select **Provisioning > Detach Replication Volume**.
2. In the main panel, click **Detach Replication Volume**. A message indicates whether the task succeeded or failed. When a volume is detached its status is shown as Detached.

Stopping a vdisk

Stopping a vdisk is part of the process for moving a secondary volume into a secondary system. The process to move a secondary volume is:

1. In the system where the secondary volume resides:
 - a. Detach the secondary volume. For details see ["Detaching a secondary volume" \(page 132\)](#).
 - b. If the secondary volume's vdisk contains other secondary volumes, detach those volumes.
 - c. Stop the secondary volume's vdisk.
 - d. If the secondary volumes' snap pools are in other vdisks, stop those vdisks.
 - e. Move the vdisks into the secondary system. This system must support the link type that the replication set is configured to use. For example, if the replication set's link type is configured to use FC links, the secondary system must have FC ports.
2. In the secondary system:
 - a. Start the snap pools' vdisks. For details see ["Starting a vdisk" \(page 133\)](#).
 - b. Start the secondary volumes' vdisks.
 - c. Reattach the secondary volumes. For details see ["Reattaching a secondary volume" \(page 134\)](#).

Before stopping a vdisk, ensure that all secondary volumes that it contains are detached. When a vdisk is stopped:

- The volumes in the vdisk become inaccessible to hosts.
- Its cached data is flushed to disk.
- Removing its disks will not cause the system to report errors or to attempt reconstruction.

NOTE: You cannot stop a vdisk that contains a primary volume.

NOTE: If a secondary volume and its snap pool are in different vdisks, you cannot stop the snap pool's vdisk until you stop the secondary volume's vdisk.

To stop a vdisk

1. In the Configuration View panel, right-click the vdisk and select **Provisioning > Stop Vdisk**.
2. In the main panel, click **Stop Vdisk**. A confirmation prompt appears.
3. Click **Yes** to continue; otherwise, click **No**. If you clicked Yes, the stop operation begins. A message indicates whether the task succeeded or failed. If the stop operation succeeds, the vdisk's health is shown as Unknown, its status is shown as STOP, and its subcomponents are no longer displayed in the Configuration View panel.
4. If the stop operation succeeded for the secondary volume's vdisk and for its snap pool's vdisk (if applicable), you can move the disks into the remote system.

Starting a vdisk

Starting a vdisk is part of the process for moving a secondary volume from a primary system into a secondary system. The process to move a secondary volume is:

1. In the system where the secondary volume resides:
 - a. Detach the secondary volume. For details see ["Detaching a secondary volume" \(page 132\)](#).
 - b. If the secondary volume's vdisk contains other secondary volumes, detach those volumes.

- c. Stop the secondary volume's vdisk. For details see ["Stopping a vdisk" \(page 133\)](#).
 - d. If the secondary volumes' snap pools are in other vdisks, stop those vdisks.
 - e. Move the vdisks into the secondary system. This system must support the link type that the replication set is configured to use. For example, if the replication set's link type is configured to use FC links, the secondary system must have FC ports.
2. In the secondary system:
 - a. Start the snap pools' vdisks.
 - b. Start the secondary volumes' vdisks.
 - c. Reattach the secondary volumes. For details see ["Reattaching a secondary volume" \(page 134\)](#).

To start a vdisk

1. In the Configuration View panel, right-click a stopped vdisk and select **Provisioning > Start Vdisk**.
2. In the main panel, click **Start Vdisk**. A message indicates whether the task succeeded or failed.

NOTE: If the replication set was deleted while the secondary volume's vdisk was stopped, restarting the vdisk will make the set partially reappear. To clean up this remnant, reattach the secondary volume, set it to be the primary volume (by using the Set Replication Primary Volume panel on [page 135](#)), and then delete the replication set again.

Reattaching a secondary volume

Reattaching a secondary volume is the last part of the process for moving a secondary volume from a primary system into a secondary system. The process to move a secondary volume is:

1. In the system where the secondary volume resides:
 - a. Detach the secondary volume. For details see ["Detaching a secondary volume" \(page 132\)](#).
 - b. If the secondary volume's vdisk contains other secondary volumes, detach those volumes.
 - c. Stop the secondary volume's vdisk. For details see ["Stopping a vdisk" \(page 133\)](#).
 - d. If the secondary volumes' snap pools are in other vdisks, stop those vdisks.
 - e. Move the vdisks into the secondary system. This system must support the link type that the replication set is configured to use. For example, if the replication set's link type is configured to use FC links, the secondary system must have FC ports.
2. In the secondary system:
 - a. Start the snap pools' vdisks. For details see ["Starting a vdisk" \(page 133\)](#).
 - b. Start the secondary volumes' vdisks.
 - c. Reattach the secondary volumes.

To reattach a secondary volume

1. In the Configuration View panel, right-click the secondary volume and select **Provisioning > Reattach Replication Volume**.
2. In the main panel, click **Reattach Replication Volume**. A message indicates whether the task succeeded or failed.
 - If the task succeeds, the secondary volume's status changes to "Establishing proxy" while it is establishing the connection to the remote (primary) system in preparation for replication; then the status changes to Online. The replication set is ready to resume replication operations.
 - If the reattach operation fails and says it is unable to get the primary volume's link type, the vdisk that contains the secondary volume may not have completed its startup activities. Wait approximately one minute for these activities to complete, then retry the operation. If this message continues to occur, check the event log to better understand the condition and for an indication of how to correct it.

NOTE: If the secondary system does not support the link type that the replication set is configured to use, the secondary volume will be attached with the wrong link type. To fix this, repeat process steps 1 and 2 above to move the secondary volume into a system that supports the required link type.

Exporting a replication image to a snapshot

If the system is licensed to use remote replication, you can export a replication image to a new standard snapshot. For example, you could export a replication image from a secondary volume for use on the remote system. The standard snapshot will reside in the same snap pool, take a snapshot license, and be independent of the primary replication image, which can continue to be used as a sync point. The standard snapshot can be used like any other standard snapshot, and changes to it will not affect the replication image.

The standard snapshot is subject to the snap pool's deletion policies. If the snap pool reaches its critical threshold, the snapshot may be deleted, even if it is mapped. If you want to preserve the standard snapshot's data, you can create a standard volume from the snapshot; see "[Creating a volume copy](#)" (page 71).

NOTE: The export task will not succeed if the resulting snapshot would exceed license limits.

To export a replication image to a snapshot

1. In the Configuration View panel, right-click a replication image and select **Provisioning > Export Snapshot**.
2. In the main panel, optionally change the default name for the snapshot. A snapshot name is case sensitive; cannot already exist in a vdisk; cannot include a comma, double quote, or backslash; and can have a maximum of 20 bytes.
3. Click **Export Snapshot**. A message specifies whether the task succeeded or failed.
4. Click **OK**. If the task succeeds, in the Configuration View panel the snapshot appears under the secondary volume on the remote system.

Changing the primary volume for a replication set

If a replication set's primary system goes offline, and the secondary volume is in a remote system, you can set the secondary volume to be the primary volume so hosts can access that volume and the replicated data it contains.

When the secondary volume becomes the primary volume, it only retains the replication images that the primary volume had and deletes any images that the primary volume did not have. Because the secondary volume may not have successfully replicated all the images associated with the primary volume, the secondary volume might have a subset of the primary volume's images.

If the primary system comes back online, you can set its volume to be the primary volume again.

The following procedures apply to a replication set in which the primary and secondary volume are in separate storage systems. To change the primary volume for a replication set in which both volumes are in the same system (local replication), use the CLI `set replication-primary-volume` command.

To change the secondary volume of a replication set to be its primary volume

1. On the secondary system, in the Configuration View panel, right-click the secondary volume and select **Provisioning > Set Replication Primary Volume**.
2. In the main panel, select the secondary volume in the list.
3. Click **Set Replication Primary Volume**. In the Configuration View panel, the volume's designation changes from Secondary Volume to Primary Volume.

NOTE: The offline primary volume remains designated a Primary Volume.

To change the primary volume of a replication set back to its original primary volume

1. On the primary system:
 - a. Create a standard snapshot ([page 69](#)) to preserve the primary volume's current data state.
 - b. Remove all mappings from the original primary volume: in the Configuration View panel, right-click the original primary volume, select **Provisioning > Explicit Mappings**, record the mappings, and remove them; then select **Provisioning > Default Mapping**, record the mapping, and remove it.
 - c. Select **Provisioning > Set Replication Primary Volume**.

- d. In the main panel, in the Primary Volume list select the primary volume that is in the secondary system.
 - e. Click **Set Replication Primary Volume**. In the Configuration View panel, the original primary volume is designated a Secondary Volume.
2. On the secondary system:
- a. Replicate the primary volume ([page 129](#)) to the secondary volume that is in the primary system to synchronize their data at the last valid common sync point. This will replicate any data changes made in the secondary volume back to the original primary volume. Let the replication operation complete.

NOTE: An administrator can mount/present/map this snapshot and the snapshot taken in [step 1](#) and compare them to verify any discrepancies.

- b. Remove all mappings from the primary volume: in the Configuration View panel, right-click the primary volume, select **Provisioning > Explicit Mappings**, record the mappings, and remove them; then select **Provisioning > Default Mapping**, record the mapping, and remove it.
3. On the primary system:
- a. In the Configuration View panel, right-click the secondary volume and select **Provisioning > Set Replication Primary Volume**.
 - b. In the main panel, in the Primary Volume list select the original primary volume.
 - c. Click **Set Replication Primary Volume**. In the Configuration View panel, the original primary volume is designated a Primary Volume. (The replication set now has a primary volume in each system.)
 - d. Re-create the mappings for the primary volume: in the Configuration View panel, right-click the primary volume, select **Provisioning > Default Mapping**, and re-create the default mapping that you recorded; then select **Provisioning > Explicit Mappings** and re-create the explicit mappings that you recorded.
4. On the secondary system:
- a. In the Configuration View panel, right-click the primary volume and select **Provisioning > Set Replication Primary Volume**.
 - b. In the main panel, in the Primary Volume list select the original primary volume.
 - c. Click **Set Replication Primary Volume**. The replication set once again has its primary volume in the primary system, and its secondary volume in the secondary system.

Viewing replication properties, addresses, and images for a volume

In the Configuration View panel, right-click a volume and select **View > Overview**. The Volume Overview table shows:

- As described in "[Viewing information about a volume](#)" ([page 104](#)): the capacity and space usage of the volume; the quantity of mappings for the volume; and the quantity of task schedules for the volume
- The quantity of replication addresses for the volume
- The quantity of replication images for the volume

For descriptions of storage-space color codes, see "[About storage-space color codes](#)" ([page 31](#)).

Select a component to see more information about it.

Replication properties

For a *local* primary or secondary volume, the Replication Properties for *Volume* table shows:

- Name. Replication volume name.
- Serial Number. Replication volume serial number.
- Status. Replication volume status:
 - Initializing: The initial (full) replication to the volume is in progress.
 - Online: The volume is online and is consistent with the last replicated image.
 - Inconsistent: The volume is online but is in an inconsistent state. A full replication is required to initialize it.
 - Replicating: The volume is online and replication is in progress.

- Replicate-delay: The volume is online but the in-progress replication has been temporarily delayed; a retry is occurring.
- Suspended: The volume is online but the in-progress replication has been suspended.
- Offline: The volume cannot be accessed or is unusable due to an error.
- Establishing proxy: The volume is establishing a proxy connection to a remote volume. This will occur when a detached secondary volume is reattached and is re-establishing a connection with the primary system in preparation for replication.
- Detached: The volume is detached for removal.
- Status-Reason. More information about the status value, or N/A for Online status.
- Monitor. Replication volume monitoring status:
 - OK: Communication to the remote volume is successfully occurring on the FC or iSCSI network.
 - Failed: Communication to the remote volume has failed because of an FC or iSCSI network issue or because the remote volume has gone offline.
- Location. Local or Remote.
- Primary Volume Name. Primary volume name. If the replication set has a primary-volume conflict, all associated primary volumes are displayed.
- Primary Volume Serial Number. Primary volume serial number. If the replication set has a primary-volume conflict, all associated primary volumes are displayed.
- Primary Volume Status. Primary volume status: Online, Offline, Conflict, or N/A.
- Maximum Number of Queued Images. Number of replication images to consider when determining the next image to replicate. Used only if the On Collision parameter is set to Oldest.
- Maximum Retry Time (Seconds). Amount of time in seconds that the replication volume should retry a replication operation on any specific image when errors occur. Used only if the On Error parameter is set to Retry.
- On Error. Error policy to invoke when errors occur during the replication process: Retry or Suspend.
- Link Type. Type of ports used to link the primary and secondary volumes: FC or iSCSI.
- On Collision. Collision policy used to determine the next image to replicate when multiple replication images are queued: Newest or Oldest.
- Monitor Interval. Interval in seconds at which the primary volume should query the secondary volume.
- Priority. Priority of the replication process on the replication volume: Low, Medium, or High.
- Connection Status.
 - Not Attempted. Communication has not been attempted to the remote volume.
 - Online. The volumes in the replication set have a valid connection but communication is not currently active.
 - Active. Communication is currently active to the remote volume.
 - Offline. No connection is available to the remote system.
- Connection Time. Date and time of the last communication with the remote volume, or N/A.

Replication addresses

If any remote port addresses are associated with this volume, when you select the Replication Addresses component, the Replication Addresses table shows:

- Connected Ports.
 - For a remote primary or secondary volume, this field shows the IDs of up to two hosts ports in the local system that are connected to the remote system. If two ports are connected but only one is shown, this indicates that a problem is preventing half the available bandwidth from being used.
 - For a local primary or secondary volume, this field shows N/A.
- Remote Address. The address of each host port in the remote system through which the volume is accessible.

Replication images

If any replication images exist for this volume, when you select the Replication Images component, the Replication Images table shows information about each image. For the selected image, the Replication Images table shows:

- Image Serial Number. Replication image serial number.
- Image Name. User-defined name assigned to the primary replication image.
- Snapshot Serial Number. Replication snapshot serial number associated with the image. The replication snapshot is associated with the replication volume specified in the request.
- Snapshot Name. Replication snapshot name associated with the image. For a secondary replication image, this value is not filled in until the replication is completed.
- Creation Date/Time. Date and time when the replication image was created on the replication volume.

Viewing information about a remote primary or secondary volume

In the Configuration View panel, right-click a *remote* primary or secondary volume and select **View > Overview**. The Replication Volume Overview table shows:

- Replication properties for the volume
- The quantity of replication addresses for the volume
- The quantity of replication images for the volume

Select a component to see more information about it.

Replication properties

When you select the Replication component a table shows the volume's replication properties, including the volume's name, serial number, status, status reason, monitor status, and location (local or remote); primary volume name, serial number, and status; maximum number of queued images, maximum retry time, error policy, collision policy, link type, monitor interval, and priority; and connection status and last connection date/time.

Replication addresses

When you select the Replication Addresses component a table shows:

- Connected Ports.
 - For a remote primary or secondary volume, this field shows the IDs of up to two hosts ports in the local system that are connected to the remote system. If two ports are connected but only one is shown, this indicates that a problem is preventing half the available bandwidth from being used.
 - For a local primary or secondary volume, this field shows N/A.
- Remote Address. The address of each host port in the remote system through which the volume is accessible.

Replication image properties

When you select the Replication Images component a table shows replication image details including the image serial number and name, snapshot serial number and name, and image creation date/time.

Viewing information about a replication image

In the Configuration View panel, right-click a replication image and select **View > Overview**. The Replication Image Overview table shows:

- Replication status properties
- Primary volume snapshot properties
- Secondary volume snapshot properties

Select a component to see more information about it.

Replication status properties

When you select the Status component a table shows the status, progress, start date/time, date/time of last update, date/time the replication was suspended, estimated completion time, elapsed or total replication time (including any suspension time). The panel also shows the replication image's serial number.

Primary volume snapshot properties

If the snapshot is on the local system, when you select the Primary Volume Snapshot component a table shows the:

- Name and serial number of the vdisk containing the snapshot
- Snapshot name, creation date/time, status, and status reason
- Primary volume name
- Snap pool name
- Amounts of total, unique, and shared data associated with the snapshot
- Default and user-specified retention priorities for this type of snapshot
- Snapshot type

If the snapshot is on a remote system, when you select the Primary Volume Snapshot component a table shows the snapshot serial number and creation date/time.

Secondary volume snapshot properties

If the snapshot is on the local system, when you select the Secondary Volume Snapshot component a table shows the:

- Name and serial number of the vdisk containing the snapshot
- Snapshot name, creation date/time, status, and status reason
- Primary volume name
- Snap pool name
- Amounts of total, unique, and shared data associated with the snapshot
- Default and user-specified retention priorities for this type of snapshot
- Snapshot type

If the snapshot is on a remote system, when you select the Secondary Volume Snapshot component a table shows the snapshot serial number and creation date/time.

A SNMP reference

This appendix describes the Simple Network Management Protocol (SNMP) capabilities that AssuredSAN 4004 Series storage systems support. This includes standard MIB-II, the FibreAlliance SNMP Management Information Base (MIB) version 2.2 objects, and enterprise traps.

AssuredSAN 4004 Series 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

AssuredSAN 4004 Series 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`) is based on the vendor name followed by “.2.” and the identifier for the particular product model. For example, the object identifier for AssuredSAN 4004 Series storage systems is 1.3.6.1.4.1.11.2.347. System uptime is an offset from the first time this object is read.

In the system group, all objects can be read. The contact, name, and location objects can be set.

In the interfaces group, an internal PPP interface is documented, but it is not reachable from external to the device.

The address translation (at) and external gateway protocol (egp) groups are not supported.

Enterprise traps

Traps can be generated in response to events 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, `dhtraps.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 RAIDar.

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). For a full description of this MIB, go to: www.emc.com/microsites/fibrealliance.

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 RAIDar. The statistics section is not implemented.

The following table lists the MIB objects, their descriptions and the value set in a AssuredSAN 4004 Series storage system. Unless specified otherwise, objects are *not* settable.

Table 13 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> fimestamp 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> fimestamp 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> fimestamp 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 13 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
connUnitControl	Not supported	invalid(2) for an SNMP GET operation and not settable through an SNMP SET operation.

Table 13 FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
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 RAIDar.	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 147)
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	Description of a component to which the revision corresponds	See "External details for connUnitRevsTable" (page 147)
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 148)
connUnitSensorName	Textual identification of the sensor intended primarily for operator use	See "External details for connUnitSensorTable" (page 148)
connUnitSensorStatus	Status indicated by the sensor	ok(3), warning(4), or failed(5) as appropriate for FRUs that are present, or other(2) if FRU is not present.
connUnitSensorInfo	Not supported	Empty string
connUnitSensorMessage	Description the sensor status as a message	connUnitSensorName followed by the appropriate sensor reading. Temperatures display in both Celsius and Fahrenheit; for example, CPU Temperature (Controller Module A): 48C 118F). Reports "Not installed" or "Offline" if data is not available.
connUnitSensorType	Type of component being monitored by this sensor	See "External details for connUnitSensorTable" (page 148)
connUnitSensorCharacteristic	Characteristics being monitored by this sensor	See "External details for connUnitSensorTable" (page 148)

Table 13 FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
<code>connUnitPortTable</code>	Includes the following objects as specified by the FA MIB2.2 Spec	
<code>connUnitPortUnitId</code>	<code>connUnitId</code> of the connectivity unit that contains this port	Same as <code>connUnitId</code>
<code>connUnitPortIndex</code>	Unique value for each <code>connUnitPortEntry</code> between 1 and <code>connUnitNumPorts</code>	Unique value for each port, between 1 and the number of ports
<code>connUnitPortType</code>	Port type	not-present(3), or n-port(5) for point-to-point topology, or l-port(6)
<code>connUnitPortFCClassCap</code>	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
<code>connUnitPortFCClassOp</code>	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
<code>connUnitPortState</code>	State of the port hardware	unknown(1), online(2), offline(3), bypassed(4)
<code>connUnitPortStatus</code>	Overall protocol status for the port	unknown(1), unused(2), ok(3), warning(4), failure(5), notparticipating(6), initializing(7), bypass(8)
<code>connUnitPortTransmitterType</code>	Technology of the port transceiver	unknown(1) for Fibre Channel ports
<code>connUnitPortModuleType</code>	Module type of the port connector	unknown(1)
<code>connUnitPortWwn</code>	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
<code>connUnitPortFCId</code>	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
<code>connUnitPortSn</code>	Serial number of the unit (for example, for a GBIC). If this is not applicable, returns an empty string.	Empty string
<code>connUnitPortRevision</code>	Port revision (for example, for a GBIC)	Empty string
<code>connUnitPortVendor</code>	Port vendor (for example, for a GBIC)	Empty string
<code>connUnitPortSpeed</code>	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
<code>connUnitPortControl</code>	Not supported	invalid(2) for an SNMP GET operation and not settable through an SNMP SET operation
<code>connUnitPortName</code>	String describing the addressed port	See "External details for connUnitPortTable" (page 149)
<code>connUnitPortPhysicalNumber</code>	Port number represented on the hardware	Port number represented on the hardware
<code>connUnitPortStatObject</code>	Not supported	0 (No statistics available)

Table 13 FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
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: <i>DDMMYYYY HHMMSS</i>	0 for logged events that occurred prior to or at startup
connUnitSEventTime	sysuptime timestamp when the event occurred	0 at startup
connUnitEventSeverity	Event severity level	error(5), warning(6) or info(8)
connUnitEventType	Type of this event	As defined in CAPI
connUnitEventObject	Not used	0
connUnitEventDescr	Text description of this event	Formatted event, including relevant parameters or values
connUnitLinkTable	Not supported	N/A
connUnitPortStatFabricTable	Not supported	N/A
connUnitPortStatSCSITable	Not supported	N/A
connUnitPortStatLANTable	Not supported	N/A
SNMP TRAPS	The following SNMP traps are supported	
trapMaxClients	Maximum number of trap clients	1
trapClientCount	Number of trap clients currently enabled	1 if traps enabled; 0 if traps not enabled
connUnitEventTrap	This trap is generated each time an event occurs that passes the connUnitEventFilter and the trapRegFilter	N/A
trapRegTable	Includes the following objects per the FA MIB2.2 Spec	
trapRegIpAddress	IP address of a client registered for traps	IP address set through Telnet
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)

Table 13 FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
trapRegRowState	Specifies the state of the row	READ: rowActive(3) if traps are enabled through Telnet; otherwise rowInactive(2) WRITE: Not supported
Enterprise-specific fields	Includes the following objects	
cpqSiSysSerialNum	System serial number	For example, 3CL8Y40991
cpqSiSysProductId	System product ID	For example, 481321-001
cpqSiProductName	System product name	For example, DH4824
cpqHoMibStatusArray	An array of MIB status structures. Octets 0–3 in block 0 are reserved for systems management and serve as an aggregate of the other MIBs.	Octet 0: 0. Octet 1 (overall status): 0 = Not available; 1 = Unknown/other; 2 = OK/normal; 3 = Degraded/warning; 4 = Failed/critical Octet 2 (system flags): 9 = device is not a server and web-based management is enabled Octet 3 (device type): 14 = enclosure For example, 00.02.09.14 (hex)
cpqHoGUID	Globally unique identifier formed from the product ID and serial number	For example, 4813213CL8Y40991

External details for certain FA MIB 2.2 objects

Tables in this section specify values for certain objects described in [Table 13](#).

External details for connUnitRevsTable

Table 14 connUnitRevsTable index and description values

connUnitRevsIndex	connUnitRevsDescription
1	Firmware revision for Storage Controller (Controller A)
2	Firmware revision for Storage Controller (Controller B)
3	Firmware revision for Memory Controller (Controller A)
4	Firmware revision for Memory Controller (Controller B)
5	Firmware revision for Storage Controller loader (Controller A)
6	Firmware revision for Storage Controller loader (Controller B)
7	Firmware revision for Management Controller (Controller A)
8	Firmware revision for Management Controller (Controller B)
9	Firmware revision for MC loader (Controller A)
10	Firmware revision for MC loader (Controller B)
11	Firmware Revision for Unified CPLD (Controller A)
12	Firmware Revision for Unified CPLD (Controller B)

Table 14 connUnitRevsTable index and description values (continued)

connUnitRevsIndex	connUnitRevsDescription
13	Firmware Revision for Expander (Controller A)
14	Firmware Revision for Expander (Controller B)
15	Hardware Revision for Controller A
16	Hardware Revision for Controller B

External details for connUnitSensorTable

Table 15 connUnitSensorTable index, name, type, and characteristic values

connUnitSensorIndex	connUnitSensorName	connUnitSensorType	connUnitSensorCharacteristic
1	CPU Temperature (Controller A)	board(8)	temperature(3)
2	CPU Temperature (Controller B)	board(8)	temperature(3)
3	FPGA Temperature (Controller A)	board(8)	temperature(3)
4	FPGA Temperature (Controller B)	board(8)	temperature(3)
5	Onboard Temperature 1 (Controller A)	board(8)	temperature(3)
6	Onboard Temperature 1 (Controller B)	board(8)	temperature(3)
7	Onboard Temperature 2 (Controller A)	board(8)	temperature(3)
8	Onboard Temperature 2 (Controller B)	board(8)	temperature(3)
9	Capacitor Temperature (Controller A)	board(8)	temperature(3)
10	Capacitor Temperature (Controller B)	board(8)	temperature(3)
11	CM Temperature (Controller A)	enclosure(7)	temperature(3)
12	CM Temperature (Controller B)	enclosure(7)	temperature(3)
13	Power Supply 1 Temperature	enclosure(7)	temperature(3)
14	Power Supply 2 Temperature	enclosure(7)	temperature(3)
15	Capacitor Pack Voltage (Controller A)	board(8)	power(9)
16	Capacitor Pack Voltage (Controller B)	board(8)	power(9)
17	Capacitor Cell 1 Voltage (Controller A)	board(8)	power(9)
18	Capacitor Cell 1 Voltage (Controller B)	board(8)	power(9)
19	Capacitor Cell 2 Voltage (Controller A)	board(8)	power(9)
20	Capacitor Cell 2 Voltage (Controller B)	board(8)	power(9)
21	Capacitor Cell 3 Voltage (Controller A)	board(8)	power(9)
22	Capacitor Cell 3 Voltage (Controller B)	board(8)	power(9)
23	Capacitor Cell 4 Voltage (Controller A)	board(8)	power(9)
24	Capacitor Cell 4 Voltage (Controller B)	board(8)	power(9)
25	Capacitor Charge Current (Controller A)	board(8)	currentValue(6)
26	Capacitor Charge Current (Controller B)	board(8)	currentValue(6)
27	Power Supply 1 Voltage, 12V	power-supply(5)	power(9)
28	Power Supply 1 Voltage, 5V	power-supply(5)	power(9)

Table 15 connUnitSensorTable index, name, type, and characteristic values (continued)

connUnitSensorIndex	connUnitSensorName	connUnitSensorType	connUnitSensorCharacteristic
29	Power Supply 1 Voltage, 3.3V	power-supply(5)	power(9)
30	Power Supply 2 Voltage, 12V	power-supply(5)	power(9)
31	Power Supply 2 Voltage, 5V	power-supply(5)	power(9)
32	Power Supply 2 Voltage, 3.3V	power-supply(5)	power(9)
33	Overall Status	enclosure(7)	other(2)

External details for connUnitPortTable

Table 16 connUnitPortTable index and name values

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

Configuring SNMP event notification in RAIDar

1. Verify that the storage system's SNMP service is enabled; see ["Changing management interface settings" \(page 42\)](#).
2. Configure and enable SNMP traps; see ["Configuring SNMP notification" \(page 44\)](#).
3. Optionally, configure a user account to receive SNMP traps; see ["Configuring user accounts" \(page 44\)](#).

SNMP management

You can manage storage devices using SNMP with a network management system such as HP OpenView, HP System Insight Manager (SIM), or HP Instant Support Enterprise Edition (ISEE). See their documentation for information about loading MIBs, configuring events, and viewing and setting group objects.

In order to view and set system group objects, SNMP must be enabled in the storage system; see ["Changing management interface settings" \(page 42\)](#). 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 ["Configuring user accounts" \(page 44\)](#). 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 enterprise traps MIB, `dhtraps.mib`. This MIB defines the content of the SNMP traps that AssuredSAN 4004 Series storage systems generate.

```
-----
-- Dot Hill Low Cost Array MIB for SNMP Traps
--
-- $Revision: 11692 $
--
-- Copyright 2005 Dot Hill Systems Corp.
-- All rights reserved. Use is subject to license terms.
--
-----

DHTRAPS-MIB
-- Last edit date: Nov 11th, 2005
DEFINITIONS ::= BEGIN
    IMPORTS
        enterprises
            FROM RFC1155-SMI
        TRAP-TYPE
            FROM RFC-1215
        connUnitEventId, connUnitEventType, connUnitEventDescr
            FROM FCMGMT-MIB;

    --Textual conventions for this MIB

-----

    -- formerly Box Hill
    dothill OBJECT IDENTIFIER ::= { enterprises 347 }

-- Related traps

dhEventInfoTrap TRAP-TYPE
    ENTERPRISE dothill
    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

dhEventWarningTrap TRAP-TYPE
    ENTERPRISE dothill
    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

dhEventErrorTrap TRAP-TYPE
    ENTERPRISE dothill
    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

dhEventCriticalTrap TRAP-TYPE
    ENTERPRISE dothill
    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

```

END

B Using FTP

Although RAIDar 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 RAIDar, prepare to use FTP:
 - a. Determine the network-port IP addresses of the system's controllers; see [Changing network interface settings](#) on [page 49](#).
 - b. Verify that the system's FTP service is enabled; see [Changing management interface settings](#) on [page 42](#).
 - c. Verify that the user you will log in as has permission to use the FTP interface; see [Modifying users](#) on [page 45](#).
2. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the destination directory for the log file.
3. Enter:

```
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" \(page 32\)](#).

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 RAIDar, prepare to use FTP:
 - a. Determine the network-port IP addresses of the system's controllers; see [Changing network interface settings on page 49](#).
 - b. Verify that the system's FTP service is enabled; see [Changing management interface settings on page 42](#).
 - c. Verify that the user you will log in as has permission to use the FTP interface; see [Modifying users on page 45](#).
2. On the log-collection system, open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the destination directory for the log file.

3. Enter:

```
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:

- *log-type* specifies the type of log data to transfer:
 - *crash1*, *crash2*, *crash3*, or *crash4*: One of the Storage Controller's four crash logs.
 - *ecdebug*: Expander Controller log.
 - *mc*: Management Controller log.
 - *scdebug*: Storage Controller log.
- *filename* is the file that will contain the transferred data. It is recommended to choose a filename that identifies the system, controller, log type, and date.

For example:

```
get managed-logs:scdebug Storage2-A_scdebug_2011_08_22.zip
```

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 RAIDar, prepare to use FTP:
 - a. Determine the network-port IP addresses of the system's controllers; see [Changing network interface settings on page 49](#).
 - b. Verify that the system's FTP service is enabled; see [Changing management interface settings on page 42](#).
 - c. Verify that the user you will log in as has permission to use the FTP interface; see [Modifying users on page 45](#).
2. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the destination directory for the log file.

3. Enter:

```
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:

- *date/time-range* is optional and specifies the time range of data to transfer, in the format: *start.yyyy-mm-dd.hh:mm. [AM|PM].end.yyyy-mm-dd.hh:mm. [AM|PM]*. The string must contain no spaces.
- *filename* is the file that will contain the data. It is recommended to choose a filename that identifies the system, controller, and date.

For example:


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

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 vdisk is quarantined, resolve the problem that is causing the vdisk to be quarantined before updating firmware. See information about events 172 and 485 in the Event Descriptions Reference Guide, and "[Removing a vdisk from quarantine](#)" (page 88).
 - 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 93](#)).
-

Updating controller-module firmware

A controller enclosure can contain one or two controller modules. In a dual-controller system, both controllers should run the same firmware version. Storage systems in a replication set must run the same firmware version. You can update the firmware in each controller module by loading a firmware file obtained from the enclosure vendor.

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

For best results, the storage system should be 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 RAIDar, 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 vdisks before starting the firmware update.
4. Restart the Management Controller (MC) in the controller to be updated; or if PFU is enabled, restart the MCs in both controllers. For the procedure, see "[Restarting or shutting down controllers](#)" (page 85).
5. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the directory containing the firmware file to load.
6. Enter:

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

For example:

```
ftp 10.1.0.9
```
7. Log in as an FTP user.

8. 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 RAIDar to perform the update, use another client, or use another FTP application.

If the Storage Controller cannot be updated, the update operation is cancelled. 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.

9. Quit the FTP session.

10. Clear your web browser's cache, then sign in to RAIDar. 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: After firmware update has completed on both controllers, if the system health is Degraded and the health reason indicates that the firmware version is incorrect, verify that you specified the correct firmware file and repeat the update. If this problem persists, contact technical support.

Updating expansion-module firmware

A drive enclosure can contain one or two expansion modules. Each expansion module contains an enclosure management processor (EMP). All modules of the same model should run the same firmware version.

Expansion-module firmware is updated in either of two ways:

- When you update controller-module firmware, all expansion modules are automatically updated to a compatible firmware version.
- You can update the firmware in each expansion module by loading a firmware file obtained from the enclosure vendor.

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

IMPORTANT: Disable PFU before attempting to update the firmware. If PFU is not disabled, it will downgrade the firmware on the corresponding expansion module. If this occurs, restart each SC.

To update expansion-module firmware

1. Obtain the appropriate firmware file and download it to your computer or network.
2. If you want to update all expansion modules, continue with the next step; otherwise, in RAIDar, determine the address of each expansion module to update:
 - a. In the Configuration View panel, select a drive enclosure.
 - b. In the enclosure properties table, note each EMP's bus ID and target ID values. For example, 0 and 63, and 1 and 63. Bus 0 is the bus that is native to a given controller, while bus 1 is an alternate path through the partner controller. It is recommended to perform update tasks consistently through one controller to avoid confusion.
3. In RAIDar, 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.
4. If the system has a single controller, stop I/O to vdisks before starting the firmware update.
5. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the directory containing the firmware file to load.

6. Enter:

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

For example:

```
ftp 10.1.0.9
```

7. Log in as an FTP user.

8. Either:

- To update all expansion modules, enter:

```
put firmware-file encl
```

- To update specific expansion modules, enter:

```
put firmware-file encl:EMP-bus-ID:EMP-target-ID
```

For example:

```
put S110R01.bin encl:1:63
```

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

It typically takes 2.5 minutes to update each EMP in a drive enclosure. Wait for a message that the code load has completed.

NOTE: If the update fails, verify that you specified the correct firmware file and try the update a second time. If it fails again, contact technical support.

9. If you are updating specific expansion modules, repeat [step 8](#) for each remaining expansion module that needs to be updated.
10. Quit the FTP session.
11. Verify that each updated expansion module has the correct firmware version.

Updating disk firmware

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

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

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

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

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 RAIDar, for each disk to update:
 - a. Determine the enclosure number and slot number of the disk.
 - b. If the disk is associated with a vdisk and is single ported, determine which controller owns the vdisk.
4. In RAIDar, 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:
 - To update all disks of the type that the firmware applies to, enter:

```
put firmware-file disk
```
 - To update specific disks, enter:

```
put firmware-file disk:enclosure-ID:slot-number
```

For example:

```
put firmware-file disk:1:11
```

△ **CAUTION:** Do not power cycle 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 RAIDar.
 - b. Power cycle all enclosures as described in your product's Setup 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 RAIDar, prepare to use FTP:
 - a. Determine the network-port IP addresses of the system's controllers; see [Changing network interface settings](#) on [page 49](#).
 - b. Verify that the system's FTP service is enabled; see [Changing management interface settings](#) on [page 42](#).
 - c. Verify that the user you will log in as has permission to use the FTP interface; see [Modifying users](#) on [page 45](#).
2. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the directory that contains the certificate files.
3. Enter:

```
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.

C Using SMI-S

This appendix provides information for network administrators who are managing the storage system from a storage management application through the Storage Management Initiative Specification (SMI-S). SMI-S is a Storage Networking Industry Association (SNIA) standard that enables interoperable management for storage networks and storage devices.

SMI-S replaces multiple disparate managed object models, protocols, and transports with a single object-oriented model for each type of component in a storage network. The specification was created by SNIA to standardize storage management solutions. SMI-S enables management applications to support storage devices from multiple vendors quickly and reliably because they are no longer proprietary. SMI-S detects and manages storage elements by type, not by vendor.

The key SMI-S components are:

- Web-based Enterprise Management (WBEM). A set of management and internet standard technologies developed to unify the management of enterprise computing environments. WBEM includes the following specifications:
 - xmlCIM: defines XML elements, conforming to DTD, which can be used to represent CIM classes and instances
 - CIM Operations over HTTP: defines a mapping of CIM operations onto HTTP; used as a transport mechanism
- Common Information Model (CIM). The data model for WBEM. Provides a common definition of management information for systems, networks, applications and services, and allows for vendor extensions. SMI-S is the interpretation of CIM for storage. It provides a consistent definition and structure of data, using object-oriented techniques. The standard language used to define elements of CIM is MOF. UML is used to create a graphical representation (using boxes and lines) of objects and relationships.
- Service Location Protocol (SLP). Enables computers and other devices to find services in a local area network without prior configuration. SLP has been designed to scale from small, unmanaged networks to large enterprise networks.

Embedded SMI-S array provider

The embedded SMI-S array provider provides an implementation of SMI-S 1.5 using `cim-xml` over HTTPS. SMI-enabled management clients such as HP SIM or HP Storage Essentials can perform storage management tasks such as monitoring, configuring or event-management. The provider supports the Array and Server profiles with additional (or supporting) subprofiles. The Server profile provides a mechanism to tell the client how to connect and use the embedded provider. The Array profile has the following supporting profiles and subprofiles:

- Array profile
- Block Services package
- Physical Package package
- Health package
- Multiple Computer System subprofile
- Masking and Mapping profile
- FC Target Ports subprofile
- SAS Target Ports subprofile
- iSCSI Target Ports subprofile
- Disk Drive Lite profile
- Extent Composition subprofile
- Storage Enclosure profile
- Fan profile
- Power Supply profile
- Sensors profile
- Access Points subprofile
- Location subprofile

- Software Inventory subprofile
- Block Server Performance subprofile
- Copy Services subprofile
- Job Control subprofile
- Storage Enclosure subprofile (if expansion enclosures are attached)
- Disk Sparing subprofile
- Object Manager Adapter subprofile
- CIM Alert and Lifecycle indications are supported.
- SLPv2 is supported.
- HTTPS using SSL encryption is supported on default port 5989. HTTP is supported on default `http` port 5988. (Both ports cannot be enabled at the same time.)

SMI-S implementation

SMI-S is implemented with the following components:

- CIM server (called a CIM Object Manager or CIMOM), which listens for WBEM requests (CIM operations over HTTP) from a CIM client, and responds.
- CIM provider, which communicates to a particular type of managed resource (for example, AssuredSAN 4004 Series storage systems), and provides the CIMOM with information about them. In theory, providers for multiple types of devices (for example, AssuredSAN 4004 Series storage systems and Brocade switches) can be plugged into the same CIMOM. However, in practice, all storage vendors provide the CIMOM and a single provider together, and they do not co-exist well with solutions from other vendors.

These components may be provided in several different ways:

- Embedded agent: The hardware device has an embedded SMI-S agent. No other installation of software is required to enable management of the device.
- SMI solution: The hardware or software ships with an agent that is installed on a host. The agent needs to connect to the device and obtain unique identifying information.

SMI-S architecture

The architecture requirements for the embedded SMI-S Array provider are to work within the Management Controller (MC) architecture, use limited disk space, use limited memory resources and be as fast as a proxy provider running on a server. The CIMOM used is the open source SFCB CIMOM.

SFCB is a lightweight CIM daemon that responds to CIM client requests and supports the standard CIM XML over `http/https` protocol. The provider is a CMPI (Common Management Protocol Interface) provider and uses this interface. To reduce the memory footprint, a third-party package called CIRCLE (www.simplewbem.org) is used. For more information on SFCB go to sblim.cvs.sourceforge.net/sblim/sfcb/README?view=markup.

About the AssuredSAN 4004 Series SMI-S provider

The GL100 provider is a SMI-S 1.5 provider which passes CTP 1.5 tests. Full provisioning is supported.

The AssuredSAN 4004 Series SMI-S provider is a full-fledged embedded provider implemented in the firmware. It provides an industry-standard WBEM-based management framework. SMI-S clients can interact with this embedded provider directly and do not need an intermediate proxy provider. The provider supports active management features such as RAID provisioning.

AssuredSAN 4004 Series CNC and SAS systems are supported. The classes for Dot Hill are `DHS_XXX`. The device namespace for Dot Hill is `/root/dhs`.

The embedded CIMOM can be configured either to listen to secure SMI-S queries from the clients on port 5989 and require credentials to be provided for all queries, or to listen to unsecure SMI-S queries from the clients on port 5988. This provider implementation complies with the SNIA SMI-S specification version 1.5.0.

NOTE: Port 5989 and port 5988 cannot be enabled at the same time.

The namespace details are given below.

- Implementation Namespace - `root/dhs`
- Interop Namespace - `root/interop`

The embedded provider set includes the following providers:

- Instance Provider
- Association Provider
- Method Provider
- Indication Provider

The embedded provider supports the following CIM operations:

- `getClass`
- `enumerateClasses`
- `enumerateClassNames`
- `getInstance`
- `enumerateInstances`
- `enumerateInstaneceNames`
- `associators`
- `associatorNames`
- `references`
- `referenceNames`
- `invokeMethod`

SMI-S profiles

SMI-S is organized around profiles, which describe objects relevant for a class of storage subsystem. SMI-S includes profiles for arrays, FC HBAs, FC switches, and tape libraries. Profiles are registered with the CIM server and advertised to clients using SLP.

Table 17 Supported SMI-S profiles

Profile/subprofile/package	Description
Array profile	Describes RAID array systems. It provides a high-level overview of the array system.
Block Services package	Defines a standard expression of existing storage capacity, the assignment of capacity to Storage Pools, and allocation of capacity to be used by external devices or applications.
Physical Package package	Models information about a storage system's physical package and optionally about internal sub-packages.
Health package	Defines the general mechanisms used in expressing health in SMI-S.
Server profile	Defines the capabilities of a CIM object manager based on the communication mechanisms that it supports.
FC Target Ports profile	Models the Fibre Channel specific aspects of a target storage system.
SAS Target Ports subprofile	Models the SAS specific aspects of a target storage system.
iSCSI Target Ports subprofile	Models the iSCSI specific aspects of a target storage system.
Access Points subprofile	Provides addresses of remote access points for management services.
Fan profile	Specializes the DMTF Fan profile by adding indications.
Power Supply profile	Specializes the DMTF Power Supply profile by adding indications.
Profile Registration profile	Models the profiles registered in the object manager and associations between registration classes and domain classes implementing the profile.
Software subprofile	Models software or firmware installed on the system.
Masking and Mapping profile	Models device mapping and masking abilities for SCSI systems.
Disk Drive Lite profile	Models disk drive devices.
Extent Composition	Provides an abstraction of how it virtualizes exposable block storage elements from the underlying Primordial storage pool.
Location subprofile	Models the location details of product and its sub-components.
Sensors profile	Specializes the DMTF Sensors profile.
Software Inventory profile	Models installed and available software and firmware.
Storage Enclosure profile	Describes an enclosure that contains storage elements (e.g., disk or tape drives) and enclosure elements (e.g., fans and power supplies).
Multiple Computer System subprofile	Models multiple systems that cooperate to present a "virtual" computer system with additional capabilities or redundancy.
Copy Services subprofile	Provides the ability to create and delete local snapshots and local volume copies (clones), and to reset the synchronization state between a snapshot and its source volume.
Job Control subprofile	Provides the ability to monitor provisioning operations, such as creating volumes and snapshots, and mapping volumes to hosts.
Disk Sparing subprofile	Provides the ability to describe the current spare disk configuration, to allocate/de-allocate spare disks, and to clear the state of unavailable disk drives.
Object Manager Adapter subprofile	Allows the client to manage the Object Manager Adapters of a SMI Agent. In particular, it can be used to turn the indication service on and off.

Block Server Performance subprofile

The implementation of the block server performance subprofile allows use of the `CIM_XXXStatisticalData` classes and their associations, and the `GetStatisticsCollection`, `CreateManifestCollection`, `AddOrModifyManifest` and `RemoveManifest` methods.

The Block Server Performance subprofile collection of statistics updates at 60-second intervals.

CIM

Supported CIM operations

SFCB provides a full set of CIM operations including `GetClass`, `ModifyClass`, `CreateClass`, `DeleteClass`, `EnumerateClasses`, `EnumerateClassNames`, `GetInstance`, `DeleteInstance`, `CreateInstance`, `ModifyInstance`, `EnumerateInstances`, `EnumerateInstanceNames`, `InvokeMethod (MethodCall)`, `ExecQuery`, `Associators`, `AssociatorNames`, `References`, `ReferenceNames`, `GetQualifier`, `SetQualifier`, `DeleteQualifier`, `EnumerateQualifiers`, `GetProperty` and `SetProperty`.

CIM Alerts

The implementation of alert indications allows a subscribing CIM client to receive events such as FC cable connects, Power Supply events, Fan events, Temperature Sensor events and Disk Drive events.

If the storage system's SMI-S interface is enabled, the system will send events as indications to SMI-S clients so that SMI-S clients can monitor system performance. For information about enabling the SMI-S interface, see "[SMI-S configuration](#)" (page 169).

The event categories in [Table 18](#) pertain to FRU assemblies and certain FRU components.

Table 18 CIM Alert indication events

FRU/Event category	Corresponding SMI-S class	Operational status values that would trigger alert conditions
Controller	DHS_Controller	Down, Not Installed, OK
Hard Disk Drive	DHS_DiskDrive	Unknown, Missing, Error, Degraded, OK
Fan	DHS_PSUFan	Error, Stopped, OK
Power Supply	DHS_PSU	Unknown, Error, Other, Stressed, Degraded, OK
Temperature Sensor	DHS_OverallTempSensor	Unknown, Error, Other, Non-Recoverable Error, Degraded, OK
Battery/Super Cap	DHS_SuperCap	Unknown, Error, OK
FC Port	DHS_FCPort	Stopped, OK
SAS Port	DHS_SASTargetPort	Stopped, OK
iSCSI Port	DHS_ISCSIEthernetPort	Stopped, OK

Life cycle indications

The SMI-S interface provides CIM life cycle indications for changes in the physical and logical devices in the storage system. The SMI-S provider supports all mandatory elements and certain optional elements in SNIA SMI-S specification version 1.5.0. CIM Query Language (CQL) and Windows Management Instrumentation Query Language (WQL) are both supported, with some limitations to the CQL indication filter. The provider supports additional life cycle indications that the Windows Server 2012 operating system requires.

Table 19 Life cycle indications

Profile or subprofile	Element description and name	WQL or CQL
Block Services	<pre>SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_StoragePool</pre> <p>Send life cycle indication when a vdisk is created or deleted.</p>	Both
Block Services	<pre>SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_StorageVolume</pre> <p>Send life cycle indication when a volume is created or deleted.</p>	Both
Block Services	<pre>SELECT * FROM CIM_InstModification WHERE SourceInstance ISA CIM_LogicalDevice</pre> <p>Send life cycle indication when disk drive (or any logical device) status changes.</p>	Both
Copy Services	<pre>SELECT * FROM CIM_InstModification WHERE SourceInstance ISA CIM_StorageSynchronized AND SourceInstance.SyncState <> PreviousInstance.SyncState</pre> <p>Send life cycle indication when the snapshot synchronization state changes.</p>	CQL
Disk Drive Life	<pre>SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_DiskDrive</pre> <p>Send life cycle indication when a disk drive is inserted or removed.</p>	Both
Fan	<pre>SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_Fan</pre> <p>Send life cycle indication when a fan is powered on or off.</p>	Both
Job Control	<pre>SELECT * FROM CIM_InstModification WHERE SourceInstance ISA CIM_ConcreteJob AND SourceInstance.OperationalStatus=17 AND SourceInstance.OperationalStatus=2</pre> <p>Send life cycle indication when a create or delete operation completes for a volume, LUN, or snapshot.</p>	WQL
Masking and Mapping	<pre>SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_AuthorizedSubject</pre> <p>Send life cycle indication when a host privilege is created or deleted.</p>	Both
Masking and Mapping	<pre>SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_ProtocolController</pre> <p>Send life cycle indication when create/delete storage hardware ID (add/remove hosts).</p>	Both
Masking and Mapping	<pre>SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_ProtocolControllerForUnit</pre> <p>Send life cycle indication when a LUN is created, deleted, or modified.</p>	Both
Multiple Computer System	<pre>SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_ComputerSystem</pre> <p>Send life cycle indication when a controller is powered on or off.</p>	Both
Multiple Computer System	<pre>SELECT * FROM CIM_InstModification WHERE SourceInstance ISA CIM_ComputerSystem AND SourceInstance.OperationalStatus <> PreviousInstance.OperationalStatus</pre> <p>Send life cycle indication when a logical component degrades or upgrades the system.</p>	WQL

Table 19 Life cycle indications (continued)

Profile or subprofile	Element description and name	WQL or CQL
Multiple Computer System	<pre>SELECT * FROM CIM_InstModification WHERE SourceInstance ISA CIM_RedundancySet AND SourceInstance.RedundancyStatus <> PreviousInstance.RedundancyStatus</pre> <p>Send life cycle indication when the controller active-active configuration changes.</p>	WQL
Target Ports	<pre>SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_FCPort</pre> <p>Send life cycle indication when a target port is created or deleted.</p>	Both
Target Ports	<pre>SELECT * FROM CIM_InstModification WHERE SourceInstance ISA CIM_FCPort AND SourceInstance.OperationalStatus <> PreviousInstance.OperationalStatus</pre> <p>Send life cycle indication when the status of a target port changes.</p>	WQL

SMI-S configuration

In the default SMI-S configuration:

- The secure SMI-S protocol is enabled, which is the recommended protocol for SMI-S.
- The SMI-S interface is enabled for the manage user.

The following table lists the CLI commands relevant to the SMI-S protocol:

Action	CLI command
Enable secure SMI-S port 5989 (and disable port 5988)	<code>set protocols smis enabled</code>
Disable secure SMI-S port 5989	<code>set protocols smis disabled</code>
Enable unsecure SMI-S port 5988 (and disable port 5989)	<code>set protocols usmis disabled</code>
Enable unsecure SMI-S port 5988	<code>set protocol usmis enabled</code>
See the current status	<code>show protocols</code>

To configure the SMI-S interface for other users:

1. Log in as manage
2. If the user does not already exist, create one using this command:
`create user level manage username`
3. Type this command:
`set user username interfaces wbi,cli,smis,ftp`

Listening for managed-logs notifications

For use with the storage system's managed logs feature, the SMI-S provider can be set up to listen for notifications that log files have filled to a point that are ready to be transferred to a log-collection system. For more information about the managed logs feature, see "[About managed logs](#)" (page 32).

To set up SMI-S to listen for managed logs notifications:

1. In the CLI, enter this command:
`set advanced-settings managed-logs enabled`
2. In an SMI-S client:
 - a. Subscribe using the `SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_LogicalFile` filter.
 - b. Subscribe using the `SELECT * FROM CIM_InstDeletion WHERE SourceInstance ISA CIM_LogicalFile` filter.

Testing SMI-S

Use an SMI-S certified client for SMI-S 1.5. HP has clients such as HP SIM and HP Storage Essentials. Other common clients are Microsoft System Center, IBM Tivoli, EMC CommandCenter and CA Unicenter. Common WBEM CLI clients are Pegasus `cimcli` and Sblim's `wbemcli`.

Testing also employs a Java Swing GUI called CIM Browser. To certify that the array provider is SMI-S 1.5 compliant, SNIA requires that the providers pass the Conformance Test Program (CTP) tests.

LUN Masking and Mapping operations

The implementation of the Masking and Mapping subprofile's extrinsic methods allows CIM clients to create LUNs by mapping volumes to logical ports. The `ExposePaths` method is fully implemented and simplifies this operation to one step. The `CreateStorageHardwareID` and `DeleteStorageHardwareID` methods allow CIM clients to create and remove hosts.

Troubleshooting

Table 20 provides solutions to common SMI-S problems.

Table 20 Troubleshooting

Problem	Cause	Solution
Unable to connect to the embedded SMI-S Array provider.	SMI-S protocol is not enabled.	Log in to the array as <code>manage</code> and type: <code>set protocol smis enabled</code>
HTTP Error (Invalid username/password or 401 Unauthorized).	User preferences are configurable for each user on the storage system.	Check that the user has access to the <code>smis</code> interface and set the user preferences to support the <code>smis</code> interface, if necessary. See "Adding users" (page 44) for instructions on how to add users. Also verify the supplied credentials.
Want to connect securely as user name <code>my_xxxx</code> .	Need to add user.	Log in to the array as <code>manage</code> . Type: <code>create user level manage my_xxxuser</code> and then type: <code>set user my_xxxuser interfaces wbi,cli,smis</code>
Unable to discover via SLP.	SLP multicast has limited range (known as hops).	Move the client closer to the array or set up a SLP DA server or using unicast requests.
Unable to determine if SMI-S is running.	Initial troubleshooting.	Install <code>wbemcli</code> on a Linux system by typing: <code>apt-get install wbemcli</code> Type: <code>wbemcli -nl -t -noverify ein 'https://manage:!manage@:5989/root/dhs:cim_computersystem'</code>

D Administering a log-collection system

A *log-collection system* receives log data that is incrementally transferred from a storage system whose 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" \(page 32\)](#).

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, SNMP traps, or SMI-S 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" \(page 154\)](#).


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: `appsrv`, `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.

Glossary

Additional Sense Code/Additional Sense Code Qualifier	See ASC/ASCQ.
Advanced Encryption Standard	See AES.
AES	Advanced Encryption Standard. A specification for the encryption of data using a symmetric-key algorithm.
array	See storage system
ASC/ASCQ	Additional Sense Code/Additional Sense Code Qualifier. Information on sense data returned by a SCSI device.
auto-write-through	See AWT.
available disk	A disk that is not being used in a vdisk, is not configured as a spare, and is not in the leftover state. It is available to be configured as a part of a vdisk or as a spare. <i>See also</i> compatible disk, dedicated spare, dynamic spare, and global spare.
AWT	Auto-write-through. A setting that specifies when the RAID controller cache mode automatically changes from write-back to write-through.
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.
chunk size	The amount of contiguous data that is written to a vdisk member before moving to the next member of the vdisk.
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.
CIM Query Language	See CQL.
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. Changing the host-port mode is also known as changing the ports' personality.
comma separated values	See CSV.
Common Information Model	See CMI.
Common Information Model Object Manager	See CIMOM.
compatible disk	A disk that can be used to replace a failed member disk of a vdisk because it both has enough capacity and is of the same type (SAS SSD, enterprise SAS, or midline SAS) as the disk that failed. <i>See also</i> available disk, dedicated spare, dynamic spare, and global spare.
complex programmable logic device	See CPLD.

Configuration Application Programming Interface	See CAPI
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: an SC processor; an MC processor; a SAS expander and an EC processor; management interfaces; cache protected by a supercapacitor pack and nonvolatile memory (CompactFlash); host, expansion, network, and service ports; and midplane connectivity. In a controller enclosure, the upper controller module is designated <i>A</i> and the lower one is designated <i>B</i> .
Converged Network Controller	See CNC.
Coordinated Universal Time	See UTC.
CPLD	Complex programmable logic device. An electronic component used to build reconfigurable digital circuits. It can replace large numbers of logic gates.
CQL	CIM Query Language.
CRC	Cyclic Redundancy Check. A mathematical algorithm that, when implemented in software or hardware, can be used to detect errors in data.
CSV	Comma separated values. A format to store tabular data in plain-text form.
Cyclic Redundancy Check	See CRC.
DAS	Direct Attach Storage. A dedicated storage device that connects directly to a host without the use of a switch.
Data Encryption Standard	See DES.
dedicated spare	A disk that is reserved for use by a specific vdisk to replace a failed disk. <i>See also</i> available disk, compatible disk, dynamic spare, and global spare.
default mapping	Host-access settings that are configured when a volume is created, and that apply to all hosts that are not explicitly mapped to that volume using different settings. <i>See also</i> explicit mapping and masking.
DES	Data Encryption Standard. An algorithm for the encryption of electronic data.
DHCP	Dynamic Host Configuration Protocol. A network configuration protocol for hosts on IP networks.
Direct Attach Storage	See DAS.
Distributed Management Task Force	See DMTF.
DMTF	Distributed Management Task Force. An industry organization that develops and maintains standards for system management.
drive enclosure	An enclosure that contains one or two expansion modules. Drive enclosures can be connected to a controller enclosure to provide additional storage capacity.
drive spin down	A power-saving feature that monitors disk activity in the storage system and spins down inactive SAS and SATA disks, based on user-selectable policies.
DRM	Disaster recovery management. Storage-system firmware features that, when the Site Replication Adapter (SRA) feature is enabled, support the use of VMware's Site Recovery Manager to automate disaster-recovery failover and failback tasks. <i>See also</i> SRA.
DSD	Drive spin down. A power-saving feature that monitors disk activity in the storage system and spins down inactive disks based on user-selectable policies.

dual-port disk	A disk that is connected to both controllers so it has two data paths, achieving fault tolerance.
Dynamic Host Configuration Protocol	See DHCP.
dynamic spare	An available compatible disk that is automatically assigned, if the dynamic spares option is enabled, to replace a failed disk in a vdisk with a fault-tolerant RAID level. See <i>also</i> available disk, compatible disk, dedicated spare, and global spare.
EC	Expander Controller. A processor (located in the SAS expander in each controller module and expansion module) that controls the SAS expander and provides SES functionality. See <i>also</i> EMP, MC, and SC.
EMP	Enclosure management processor. An EC subsystem that provides SES data such as temperature, power supply and fan status, and the presence or absence of disks.
enclosure	A physical storage device that contains I/O modules, disk drives, and other FRUs.
enclosure management processor	See EMP.
Expander Controller	See EC.
expansion enclosure	See drive enclosure.
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. In a drive enclosure, the upper expansion module is designated <i>A</i> and the lower one is designated <i>B</i> .
explicit mapping	Access settings for a host to a volume that override the volume's default mapping. See <i>also</i> default mapping and masking.
extrinsic methods	Methods which are particular to a specific class in SMI-S. See <i>also</i> intrinsic methods.
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 volumes, cache data, host ID information, and LUNs and WWNs. See recovery.
FC	Fibre Channel interface protocol.
FC-AL	Fibre Channel Arbitrated Loop. The FC topology in which devices are connected in a one-way loop.
Fibre Channel Arbitrated Loop	See FC-AL.
field-programmable gate array	See FPGA.
field-replaceable unit	See FRU.
FPGA	Field-programmable gate array. An integrated circuit designed to be configured after manufacturing.
FRU	Field-replaceable unit. A part that can be removed and replaced by the user or support technician without having to send the product to a repair facility.
global spare	A compatible disk that is reserved for use by any vdisk with a fault-tolerant RAID level to replace a failed disk. See <i>also</i> available disk, compatible disk, dedicated spare, dynamic spare.
HBA	Host bus adapter. A device that facilitates I/O processing and physical connectivity between a host and the storage system.
host	An external port that the storage system is attached to. The external port may be a port in an I/O adapter in a server, or a port in a network switch. Product interfaces use the terms host and initiator interchangeably.
host port	A port on a controller module that interfaces to a host computer, either directly or through a network switch.

host bus adapter	See HBA.
image ID	A globally unique serial number that identifies the point-in-time image source for a volume. All volumes that have identical image IDs have identical data content, whether they be snapshots or stand-alone volumes.
initiator	See host.
I/O Manager	A MIB-specific term for a controller module.
I/O module	See IOM.
intrinsic methods	Methods inherited from CIM and present in all classes such as <code>getClass</code> , <code>createInstance</code> , <code>enumerateInstances</code> , and <code>associatorNames</code> in SMI-S.
IOM	Input/output module. An IOM can be either a controller module or an expansion module.
IQN	iSCSI Qualified Name.
iSCSI	Internet SCSI interface protocol.
iSNS	Internet Storage Name Service.
JBOD	"Just a bunch of disks." See drive enclosure.
jumbo frame	In an iSCSI network, a frame that can contain 9000 bytes for large data transfers. A normal frame can contain 1500 bytes.
large form factor	See LFF.
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 vdisk because the timestamp in the disk's metadata is older than the timestamp of other disks in the vdisk, or because the disk was not detected during a rescan. A leftover disk cannot be used in another vdisk 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. A type of disk drive.
LIP	Loop Initialization Primitive. An FC primitive used to determine the loop ID for a controller.
Logical Block Address	See LBA.
Logical Unit Number	See LUN.
loop	See FC-AL.
Loop Initialization Primitive	See LIP.
LUN	Logical Unit Number. A number that identifies a mapped volume to a host.
MAC address	Media Access Control Address. A unique identifier assigned to network interfaces for communication on a network.
Management Controller	See MC.
Management Information Base	See MIB.
map/mapping	Settings that specify whether a volume is presented as a storage device to a host, and how the host can access the volume. Mapping settings include an access type (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. See <i>also</i> default mapping, explicit mapping, and masking.
masking	A volume-mapping setting that specifies no access to that volume by hosts. See <i>also</i> default mapping and explicit mapping.
master volume	A volume that is enabled for snapshots and has an associated snap pool.
MC	Management Controller. The processor (located in a controller module) that is responsible for human-computer interface and computer-computer interfaces, including the WBI, CLI, and FTP interfaces, and interacts with the Storage Controller. See <i>also</i> EC and SC.

Media Access Control Address	See MAC address.
metadata	Data in the first sectors of a disk drive that stores all disk-, vdisk-, and volume-specific information including vdisk membership or spare identification, vdisk ownership, volumes and snapshots in the vdisk, host mapping of volumes, and results of the last media scrub.
MIB	Management Information Base. A database used for managing the entities in SNMP.
mount	To enable access to a volume from a host OS. See <i>also</i> host, map/mapping, and volume.
network port	The Ethernet port on a controller module through which its Management Controller is connected to the network.
network time protocol	See NTP.
NTP	Network time protocol.
object identifier	See OID.
OID	Object Identifier. In SNMP, an identifier for an object in a MIB.
orphan data	See unwritable cache data.
Partner Firmware Update	See PFU.
PFU	Partner firmware update. The automatic update of the partner controller when the user updates firmware on one controller.
PHY	One of two hardware components that form a physical connection between devices in a SAS network that enables transmission of data.
physical layer	See PHY.
point-to-point	Fibre Channel Point-to-Point topology, where two ports are directly connected.
POST	Power-on self test. Tests that run immediately after a device is powered on.
Power-On Self Test	See POST.
power supply unit	See PSU.
primary volume	The volume that is the source of data in a replication set and that can be mapped to hosts. For disaster recovery purposes, if the primary volume goes offline, a secondary volume can be designated as the primary volume. The primary volume exists in a primary vdisk in the primary (or local) storage system.
proxy volume	A virtual volume in the local system that represents a volume in a remote system. Proxy volumes are used internally by the controllers to perform actions such as transferring replication data.
PSU	Power supply unit. The power supply FRU.
recovery	In an active-active configuration, recovery is the act of returning ownership of controller resources to a controller (which was offline) from its partner controller. The resources include volumes, cache data, host ID information, and LUNs and WWNs. See <i>also</i> failover.
remote replication	Asynchronous (batch) replication of block-level data from a volume in a primary system to a volume in one or more secondary systems by creating a replication snapshot of the primary volume and copying the snapshot data to the secondary systems via Fibre Channel or iSCSI links. The capability to perform remote replication is a licensed feature (AssuredRemote).
remote syslog support	See syslog.
replication image	A conceptual term for replication snapshots that have the same image ID in primary and secondary systems. These synchronized snapshots contain identical data and can be used for disaster recovery.
replication set	Associated primary and secondary volumes that are enabled for replication and that typically reside in two physically or geographically separate storage systems. See primary volume and secondary volume.

replication snapshot	A special type of snapshot, created by the remote replication feature, that preserves the state of data of a replication set's primary volume as it existed when the snapshot was created. For a primary volume, the replication process creates a replication snapshot on both the primary system and, when the replication of primary-volume data to the secondary volume is complete, on the secondary system. Replication snapshots are unmappable and are not counted toward a license limit, although they are counted toward the system's maximum number of volumes. A replication snapshot can be exported to a regular, licensed snapshot. See <i>also</i> replication sync point.
replication sync point	The state of a replication snapshot whose corresponding primary or secondary snapshot exists and contains identical data. For a replication set, four types of sync point are identified: the only replication snapshot that is copy-complete on any secondary system is the "only sync point"; the latest replication snapshot that is copy-complete on any secondary system is the "current sync point"; the latest replication snapshot that is copy-complete on all secondary systems is the "common sync point"; a common sync point that has been superseded by a new common sync point is an "old common sync point."
SAS	Serial Attached SCSI interface protocol or disk-drive architecture.
SC	Storage Controller. The processor (located in a controller module) that is responsible for RAID controller functions. The SC is also referred to as the RAID controller. See <i>also</i> EC and MC.
SCSI Enclosure Services	See SES.
secondary volume	The volume that is the destination for data in a replication set and that is not accessible to hosts. For disaster recovery purposes, if the primary volume goes offline, a secondary volume can be designated as the primary volume. The secondary volume exists in a secondary vdisk in a secondary (or remote) storage system. The contents of a secondary volume are in a constant state of flux and are not in a consistent state while a replication is in process. Only snapshots that are associated with a secondary volume are data consistent.
secret	For use with CHAP, a password that is shared between an initiator and a target to enable authentication.
secure hash algorithm	See SHA.
secure shell	See SSH.
Secure Sockets Layer	See SSL.
SEEPROM	Serial electrically erasable programmable ROM. A type of nonvolatile (persistent if power removed) computer memory used as FRU ID devices.
Self-Monitoring Analysis and Reporting Technology	See SMART.
serial electrically erasable programmable ROM	See SEEPROM.
Service Location Protocol	See SLP.
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. A cryptographic hash function.
SLP	Service Location Protocol. Enables computers and other devices to find services in a local area network without prior configuration.
Small Footprint CIM Broker	See SFCB.

small form factor	See SFF.
SMART	Self-Monitoring Analysis and Reporting Technology. A monitoring system for disk drives that monitors reliability indicators for the purpose of anticipating disk failures and reporting those potential failures.
SMI-S	Storage Management Initiative - Specification. The SNIA standard that enables interoperable management of storage networks and storage devices. The interpretation of CIM for storage. It provides a consistent definition and structure of data, using object-oriented techniques.
snap pool	A volume that stores data that is specific to snapshots of an associated master volume, including copy-on-write data and data written explicitly to the snapshots. A snap pool cannot be mapped.
snapshot	A “virtual” volume that preserves the state of a master volume’s data as it existed when the snapshot was created. Data associated with a snapshot is recorded in both the master volume and in its associated snap pool. A snapshot can be mapped and written to. The capability to create snapshots is a licensed feature (AssuredSnap). Snapshots that can be mapped to hosts are counted against the snapshot-license limit, whereas transient and unmappable snapshots are not.
SNIA	Storage Networking Industry Association. An association regarding storage networking technology and applications.
SRA	Storage Replication Adapter. A host-based software component that allows VMware’s Site Recovery Manager to manage the storage-system firmware’s disaster recovery management (DRM) features, automating disaster-recovery failover and failback tasks. The SRA uses the CLI XML API to control the storage system. <i>See also</i> DRM.
SSD	Solid-state drive.
SSH	Secure Shell. A network protocol for secure data communication.
SSL	Secure Sockets Layer. A cryptographic protocol that provides security over the internet.
Storage Controller	See SC.
Storage Management Initiative - Specification	See SMI-S.
Storage Networking Industry Association	See SNIA.
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	Remote syslog support. A configuration that, when enabled, sends selected event messages to the syslog on a remote system.
UCS Transformation Format - 8-bit	See UTF-8.
ULP	Unified LUN Presentation. A RAID controller feature that enables a host to access mapped volumes through any controller host port. ULP incorporates Asymmetric Logical Unit Access (ALUA) extensions.
Unified LUN Presentation	See ULP.
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 and disk. Unwritable cache is also called orphan data.
UTC	Coordinated Universal Time. The primary time standard by which the world regulates clocks and time. It replaces Greenwich Mean 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 WBI interfaces.

vdisk	A virtual disk comprising the capacity of one or more disks. The number of disks that a vdisk can contain is determined by its RAID level.
virtual disk	See vdisk.
volume	A portion of the capacity of a vdisk that can be presented as a storage device to a host.
volume copy	An independent copy of the data in a volume. The capability to create volume copies is a licensed feature (AssuredCopy) that makes use of snapshot functionality.
WBEM	Web-Based Enterprise Management. A set of management and internet standard technologies developed to unify the management of enterprise computing environments.
web-based interface/web-browser interface	See WBI.
WBI	Web-based interface/web-browser interface. The primary interface for managing the system. A user can enable the use of HTTP, HTTPS for increased security, or both.
Web-Based Enterprise Management	See WBEM.
Windows Management Instrumentation Query Language	See WQL.
World Wide Name	See WWN.
World Wide Node Name	See WWNN.
World Wide Port Name	See WWPN.
WQL	Windows Management Instrumentation Query Language.
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 16

A

activity progress interface 81

ALUA 22

array

See system

asterisk (*) in option name 16

audience, document 13

B

base for size representations 30

bytes versus characters 30

C

cache

configuring auto-write-through triggers and behaviors 53

configuring host access to 52

configuring independent cache performance mode 52

configuring system settings 52

configuring volume settings 58

certificate

using FTP to install a security 160

CHAP

adding or modifying records 76

configuring 40, 48

configuring for iSCSI hosts 76

deleting records 76

overview 21

characters versus bytes 30

color codes for storage space 31

CompactFlash properties 117

configuration

browser 15

first-time 15

system limits 97

Configuration View component icons 31

Configuration View panel, using 16

Configuration Wizard, using 37

controller module properties 114

controllers

restarting or shutting down 85

using FTP to update firmware 156

using WBI to update firmware 79

conventions, document 14

Coordinated Universal Time (UTC) 30

current owner 57

D

date and time

about 30

configuring 47

debug data

saving to a file 83

debug logs

downloading 153

dedicated spare 19

dedicated spares

adding and removing 56

default mapping 22

DHCP

configuring 49

configuring with Configuration Wizard 37

disk channels

rescanning 84

disk metadata

clearing 84

disk performance

about monitoring historical 34

resetting (clearing) historical statistics 90

saving (downloading) historical statistics 91

disk properties 94, 102, 111

disk settings

configuring 50

disk state (how used) values 103

disks

configuring background scrub 54

configuring SMART 50

configuring spin down for available and global-spare 50

enabling vdisk reconstruction by replacing failed 31

maximum number 97

scheduling spin down for all 51

showing data transfer rate 112

using FTP to retrieve performance statistics 155

using FTP to update firmware 159

using WBI to update firmware 81

document

audience 13

conventions 14

prerequisite knowledge 13

related documentation 13

drive spin down

configuring for a vdisk 57

configuring for available and global-spare disks 50

scheduling for all disks 51

dynamic spare 19

dynamic spares

configuring 50

E

- EMP polling rate
 - configuring 51
- enclosure
 - viewing information about 110
- enclosure properties 94, 111
- event log
 - viewing 98
- event notification
 - configuring email settings 43
 - configuring SNMP settings 44
 - configuring with Configuration Wizard 39
 - sending a test message 86
- event severity icons 98
- expansion module properties 118
- expansion port properties 117
- explicit mapping 22

F

- firmware
 - about updating 34
 - using FTP to update controller module 156
 - using FTP to update disk drive 159
 - using FTP to update expansion module 157
 - using WBI to update controller module 79
 - using WBI to update disk 81
 - using WBI to update expansion module 80
- firmware update, monitoring progress of 81
- firmware update, partner
 - configuring 54
- firmware versions 97
- FTP
 - downloading system logs 153
 - retrieving disk-performance statistics 155
 - updating controller module firmware 156
 - updating disk drive firmware 159
 - updating expansion module firmware 157
 - using to install a security certificate 160
 - using with the log-management feature 154

G

- global spare 19
- global spares
 - adding and removing 62

H

- hardware versions 97
- help
 - using online 17
- host
 - adding 74
 - changing mappings 75
 - changing name 75
 - viewing information about 110
- host access to cache
 - configuring 52
- host mapping properties 110

- host port properties
 - FC 115
 - iSCSI 116
 - SAS 117
- host ports
 - checking links in local system 90
 - checking links to remote system 90
 - configuring 47
 - configuring with Configuration Wizard 40
 - maximum number 97
 - resetting 83
- host properties 110
- hosts
 - about 21
 - removing 75
 - viewing information about all 109

I

- I/O module properties 118
- icons
 - event severity 98
 - storage-system component 31
 - WBI communication status 16
- In port properties 118
- independent cache performance mode
 - configuring 52
- iSCSI host security 21
- iSCSI IP version
 - configuring 40, 48
- iSNS
 - configuring 40, 48

J

- jumbo frames
 - configuring 40, 48

L

- leftover disk 84
- licensed features
 - creating a temporary license 41
 - installing a permanent license 41
 - remote replication 121
 - snapshot limit 24
 - Storage Replication Adapter (SRA) 28
 - using FTP to install license file 160
 - viewing status 97
 - Virtual Disk Service (VDS) hardware provider 28
 - volume copy 26
 - Volume Shadow Copy Service (VSS) hardware provider 28
- link rate adjustment 112
- link speed
 - configuring FC 40, 48
- links
 - checking between controllers in local system 90
 - checking between local and remote systems 90
- log data
 - saving to a file 83

- log management
 - about [32](#)
 - sending a test message [86](#)
 - using FTP [154](#)
- log-collection system
 - administering [171](#)
- logs
 - downloading debug [153](#)
- LUNs
 - configuring response to missing [52](#)
 - maximum number [97](#)
- M**
- managed logs
 - about [32](#)
 - administering a log-collection system [171](#)
 - enabling/disabling [55](#)
- management interface services
 - configuring [42](#)
 - configuring with Configuration Wizard [38](#)
- mapping volumes
 - See volume mapping
- masked volume [22](#)
- master volumes
 - about [24](#)
- maximum physical and logical entities supported [97](#)
- metadata
 - clearing disk [84](#)
- MIB
 - See SNMP
- missing LUN response
 - configuring [52](#)
- modified snapshot data, deleting
 - about [25](#)
- N**
- network port [37](#)
- network port properties [115](#)
- network ports
 - configuring [49](#)
 - configuring with Configuration Wizard [37](#)
- NTP
 - about [30](#)
 - configuring [47](#)
- O**
- Out port properties [117](#), [118](#)
- P**
- partner firmware update
 - configuring [54](#)
- passwords
 - See users
- performance monitoring
 - See disk performance
- policies and thresholds, snap-pool [108](#)
- power supply properties [114](#)
- preferred owner [57](#)
- prerequisite knowledge, document [13](#)
- primary volume
 - changing for a replication set [135](#)
- priority
 - configuring utility [55](#)
- Provisioning Wizard
 - using to create a vdisk with volumes and mappings [59](#)
- provisioning, first-time [15](#)
- R**
- RAID levels
 - about [28](#)
- RAIDar Storage Management Utility
 - See WBI
- read-ahead caching
 - optimizing [23](#)
- related documentation [13](#)
- remote replication
 - about [121](#)
- remote systems
 - about managing [24](#)
 - adding [55](#)
 - checking links from local system [90](#)
 - removing [55](#)
 - viewing information about [119](#)
- replication disaster recovery [135](#)
- replication image
 - viewing information about [138](#)
- replication image primary-snapshot properties [139](#)
- replication image replication-status properties [139](#)
- replication image secondary-snapshot properties [139](#)
- replication set
 - changing the primary volume [135](#)
 - detaching a secondary volume [132](#)
 - reattaching a secondary volume [134](#)
- Replication Setup Wizard
 - using to set up replication for a volume or snapshot [127](#)
- replication snapshot size [122](#)
- replication volume
 - viewing information about a remote primary or secondary [138](#)
- rescan disk channels [84](#)
- restarting controllers [85](#)
- restoring the system's default configuration settings [84](#)
- reverting volume data
 - See rolling back volume data
- rolling back volume data
 - about [25](#)
- S**
- schedule
 - deleting [78](#)
 - modifying [77](#)
- schedule properties [96](#), [108](#)

- scheduling
 - snapshot [69](#)
 - snapshot reset [70](#)
 - volume copy [71](#)
- scrub
 - configuring background disk [54](#)
 - configuring background vdisk [54](#)
- SCSI MODE SELECT command
 - configuring handling of [52](#)
- SCSI SYNCHRONIZE CACHE command
 - configuring handling of [52](#)
- secondary volume
 - detaching [132](#)
 - reattaching [134](#)
- security certificate
 - using FTP to install [160](#)
- selective storage presentation
 - See volume mapping
- shared data (snapshot) [107](#)
- shutting down controllers [85](#)
- sign out, auto
 - setting user [45, 46](#)
 - viewing remaining time [16](#)
- signing in to the WBI [15](#)
- signing out of the WBI [16](#)
- single-controller system data-protection tips [32](#)
- size representations
 - about [30](#)
 - replication snapshot [122](#)
- SMART
 - configuring [50](#)
- snap data [107](#)
- snap pool
 - about [24](#)
 - creating [74](#)
 - expanding [89](#)
 - viewing information about [108](#)
- snap pool properties [96, 103, 108](#)
- snap pool thresholds and policies [108](#)
- snap pools
 - deleting [74](#)
 - renaming [58](#)
- snapshot
 - creating [69](#)
 - exporting replication image to [135](#)
 - preparing replication by using the Replication Setup Wizard [127](#)
 - replicating [130](#)
 - resetting to current data in master volume [70](#)
 - viewing information about [106](#)
- snapshot mapping properties [107](#)
- snapshot properties [96, 106, 109](#)
- snapshots
 - about [24](#)
 - creating for multiple volumes [68](#)
 - deleting [70](#)
 - renaming [58](#)

- SNMP
 - configuring traps [149](#)
 - enterprise trap MIB [150](#)
 - enterprise traps [141](#)
 - external details for connUnitPortTable [149](#)
 - external details for connUnitRevsTable [147](#)
 - external details for connUnitSensorTable [148](#)
 - FA MIB 2.2 behavior [142](#)
 - FA MIB 2.2 objects, descriptions, and values [142](#)
 - management [149](#)
 - MIB-II behavior [141](#)
 - overview [141](#)
 - setting event notification [149](#)
- sorting a table [16](#)
- spares
 - about [19](#)
 - See also dedicated spare, dynamic spare, and global spare
- SRA
 - about [28](#)
- Storage Replication Adapter (SRA)
 - See SRA
- storage system
 - See system
- synchronize-cache mode
 - configuring [52](#)
- syslog
 - sending a test message [86](#)
- system
 - configuration limits [97](#)
 - data-protection tips for a single-controller [32](#)
 - restoring default configuration settings [84](#)
 - viewing event log [98](#)
 - viewing information about [93](#)
- system information
 - configuring [50](#)
 - configuring with Configuration Wizard [39](#)
- system properties [93](#)
- System Status panel
 - using [16](#)
- system utilities
 - configuring [54](#)

T

- tables sorting [16](#)
- task schedule
 - See schedule
- temperature
 - configuring controller shutdown for high [53](#)
- thresholds and policies, snap-pool [108](#)
- time and date
 - about [30](#)
 - configuring [47](#)

U

ULP 22

unique data (snapshot) 107

units for size representations 30

users

about user accounts 17

adding 44

changing default passwords with Configuration Wizard 37

maximum that can sign in 16

modifying 45

removing 47

utility priority

configuring 55

V

vdisk

aborting scrub 88

aborting verification 87

changing name 56

changing owner 57

configuring 56

configuring drive spin down 57

creating 61

creating with the Provisioning Wizard 59

expanding 86

removing from quarantine 88

scrubbing 87

starting a stopped 133

stopping 133

verifying redundant 87

viewing information about 100

vdisk health values 99, 100

vdisk performance graphs 101

vdisk properties 95, 100

vdisk reconstruction

replacing failed disks to enable 31

setting spares to enable 19

vdisk status values 99, 101

vdisks

about 18

configuring background scrub 54

deleting 62

maximum number 97

viewing information about all 99

VDS and VSS providers

about 28

volume

aborting copy 73

aborting replication 132

changing default mapping 66

changing explicit mappings 67

changing name 58

configuring 58

configuring cache settings 58

creating 63

creating a copy 71

expanding 68

preparing replication by using the Replication Setup Wizard 127

removing replication from 131

replicating 129

resuming replication 132

rolling back data 73

suspending replication 131

viewing information about 104

volume cache options

about 22

volume copy

about 26

volume mapping

about 22

changing default mapping for multiple volumes 64

changing explicit mapping for multiple volumes 65

unmapping multiple volumes 68

volume mapping properties 105

volume masking 22

volume properties 96, 103, 104, 109

volume replication addresses 137, 138

volume replication images 138

volume replication properties 136, 138

volume replication-image properties 138

volume schedule properties 105

volume set

creating 63

volumes

about 20

deleting 64

maximum number 97

W

WBI

about 15

signing in 15

signing out 16

WBI communication status icon 16

WBI session hang 16

web-browser buttons to avoid 16

web-browser interface

See WBI

web-browser setup 15

write-back caching 23

write-through caching 23

