AssuredSAN 4000 Series
RAIDar User Guide

P/N 83-00005253-10-01
Revision B
October 2012

Contents

Tables

About this guide

This guide provides information about managing an AssuredSAN™ 4000 Series storage system by using its web interface, RAIDar Storage Management Utility.

Intended audience

This guide is intended for storage system administrators.

Prerequisites

Prerequisites for using this product include knowledge of:

- Network administration
- Storage system configuration
- SAN management and DAS
- FC, SAS, and Ethernet protocols

Related documentation

Table 1 Related Documentation

For information about	See
Enhancements, known issues, and late-breaking information not included in product documentation	AssuredSAN 4000 Series 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*
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 4000 Series Setup Guide
Using the CLI to configure and manage the product	AssuredSAN 4000 Series CLI Reference Guide
Event codes and recommended actions Identifying and installing or replacing FRUs	AssuredSAN 4000 Series Service Guide

* Printed document included in product shipkit.

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

Document conventions and symbols

Table 2 Document conventions

Convention	Element
Blue text	Cross-reference links and e-mail addresses
Blue, underlined text	Web site addresses
Bold font	• 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 checkboxes
Italics font	Text emphasis
`Monospace font`	• File and directory names • System output • Code • Text typed at the command-line
`Monospace, italic font`	• 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. 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 WBI.

NOTE: Open RAIDar sessions can negatively impact system performance.

Configuring and provisioning a new storage system

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

1. Configure your web browser for RAIDar and sign in, as described in Browser setup and Signing in.
2. Set the system date and time, as described in Changing the system date and time on page 36.
3. Use the Configuration Wizard to configure other system settings, as described in Using the Configuration Wizard on page 27.
4. Use the Provisioning Wizard to create a vdisk containing storage volumes, and optionally to map the volumes to hosts, as described in Using the Provisioning Wizard on page 47.
5. 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.
6. Verify that controller modules and expansion modules have the latest firmware, as described in Viewing information about the system on page 71 and Updating firmware on page 59.

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

Browser setup

- Use Mozilla Firefox 3 or later, or Microsoft Internet Explorer 8 or later.
- 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. For Internet Explorer 8, adding each controller's network IP address as a trusted site can avoid access issues.
  - Verify that the browser is set to allow cookies at least for the IP addresses of the storage-system network ports.
- To see the help window in Microsoft Internet Explorer, you must enable pop-up windows.

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.**0**33.
- 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. Internet Explorer can run separate RAIDar sessions if you select **File > New Session**. If you do not select a new session, all instances of Internet Explorer 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 displays System Time and System Events. System Events shows 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 on page 75.
- 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 MC, and the SC, as described in the following table.

Table 3 RAIDar communication status icons

Icon	Meaning
	RAIDar can communicate with the MC, which can communicate with the SC.
	RAIDar *cannot* communicate with the MC.
	RAIDar can communicate with the MC, which *cannot* communicate with the SC.

- 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 on page 22.

- Icons shown in the Configuration View panel are described in About Configuration View icons on page 23.

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 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 on page 32. For information about the MIB, see SNMP reference on page 89.

General user accounts have these options:

- User Name.
- Password.
- User Roles. Either Monitor or Manage. You cannot change the roles of user `manage`.
  - Monitor. A user with role Monitor may view system settings but cannot change them.
  - Manage. A user with role Manage may 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.
- SMI-S Access. Allows access to the 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, or 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. 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 authentication; or none (no authentication). Authentication uses the password set with the Password option.
- Privacy Type. Either: DES; AES; or none (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 4 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 on page 33

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 36 vdisks per controller and 72 per system can exist.

A vdisk can contain different models of disks, and disks with different capacities. For example, a vdisk can include a 500 GB disk and a 750 GB disk. 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, if a RAID0 vdisk contains one 500 GB disk and four 750 GB disks, the capacity of the vdisk is equivalent to approximately five 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.

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.

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 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 on page 19
- About spares on page 15
- About volumes on page 16
- Vdisk topics in Provisioning the system on page 47
- Configuring an existing vdisk on page 43
- Verifying a vdisk on page 67
- Scrubbing a vdisk on page 67
- Viewing information about a vdisk (page 76), all vdisks (page 76), or the system (page 71)
- Removing a vdisk from quarantine on page 68

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.

Related topics

- Configuring dynamic spares on page 38
- Managing dedicated spares on page 43
- Managing global spares on page 50
- Using the Provisioning Wizard on page 47
- Removing hosts on page 56
- Viewing information about a vdisk (page 76) or all vdisks (page 76)

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 vdisk can have a maximum of 128 volumes.

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 RAID5 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 RAID5 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.

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 on page 14
- About volume mapping on page 17
- About volume cache options on page 18
- Volume topics in Provisioning the system on page 47

- Changing a volume's name on page 45
- Changing a volume's cache settings on page 45
- Viewing information about a volume on page 79

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, or a port in a network switch.

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. 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.

Related topics

- Using the Configuration Wizard on page 27
- Changing host interface settings on page 37
- Removing hosts on page 56
- Changing a host's name on page 56
- Changing host mappings on page 56
- Viewing information about a host (page 81) or all hosts (page 80)

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 example, if LUN 5 is the default for Volume1, no other volume in the storage system can use LUN 5 as its default LUN. For explicit mappings, the rules differ: LUNs used in default mappings can be reused in explicit mappings for other volumes and other hosts.

TIP: When an explicit mapping is deleted, the volume's default mapping takes effect. Therefore, it is recommended to use the same LUN for explicit mappings as for the default mapping.

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 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 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 on page 47
- Changing a volume's default mapping (page 53) or explicit mappings (page 54)
- Changing host mappings on page 56
- Viewing information about a volume (page 79), host (page 81), or all hosts (page 80)

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. 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 is triggered by two back-to-back accesses to consecutive LBA ranges, whether forward (increasing LBAs) or reverse (decreasing LBAs).

You can change the amount of data read in advance after two back-to-back reads are made. 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.

- The Default option works well for most applications: it sets one chunk for the first access in a sequential read and one stripe for all subsequent accesses. The size of the chunk is based on the chunk size used when you created the vdisk (the default is 64 KB). Non-RAID and RAID1 vdisks are considered to have a stripe size of 64 KB.
- Specific size options let you select an amount of data for all accesses.
- The Maximum option lets the controller dynamically calculate the maximum read-ahead cache size for the volume. For example, if a single volume exists, this setting enables the controller to use nearly half the memory for read-ahead cache.

  Only use Maximum when disk latencies must be absorbed by cache. For example, for read-intensive applications, you will want data that is most often read to be in cache so that the response to the read request is very fast; otherwise, the controller has to locate which disks the data is on, move it up to cache, and then send it to the host. Do not use Maximum if more than two volumes are owned by the controller on which the read-ahead setting is being made. If there are more than two volumes, there is contention on the cache as to which volume's read data should be held and which has the priority; each volume constantly overwrites the other volume's data in cache, which could result in taking a lot of the controller's processing power.
- The Disabled option turns off read-ahead cache. This is useful if the host is triggering read ahead for what are random accesses. This can happen if the host breaks up the random I/O into two smaller reads, triggering read ahead.

You can also change the optimization mode.

- The Standard read-ahead caching mode works well for typical applications where accesses are a combination of sequential and random. For example, use this mode for transaction-based and database update applications that write small files in random order.
- You can use No-mirror mode. When this mode is enabled, 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.

Related topics

- Changing a volume's cache settings on page 45
- Changing system cache settings on page 40
- Viewing information about a volume on page 79

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.

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

Table 5, Table 6, and Table 7:

- Provide examples of appropriate RAID levels for different applications
- Compare the features of different RAID levels
- Describe the expansion capability for different RAID levels

Table 5 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 6 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 RAID0 or RAID1

Table 6 RAID level comparison (continued)

RAID level	Min. disks	Description	Strengths	Weaknesses
6	4	Block-level data striping with double distributed parity	Best suited for large sequential workloads; non-sequential read and sequential read/write performance is comparable to RAID5; protects against dual disk failure	Higher redundancy cost than RAID5 because the parity overhead is twice that of RAID5; not well-suited for transaction-oriented network applications; non-sequential write performance is slower than RAID5
10 (1+0)	4	Stripes data across multiple RAID1 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 RAID5 sub-vdisks	Better random read and write performance and data protection than RAID5; supports more disks than RAID5; protects against multiple disk failures	Lower storage capacity than RAID5

Table 7 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 8 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 in Table 9.

Table 9 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 on page 13

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 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 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 on page 36

About storage-space color codes

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

Table 10 Storage-space color codes

Area	Color	Meaning
Overview panels		Total space
		Available/free space
		Used space
		Reserved/overhead space, used for parity, 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 11 Configuration View icons

Icon	Meaning	Icon	Meaning
	Show all subcomponents		Enclosure
	Hide all subcomponents		Host/initiator
	Show the component's subcomponents		Vdisk
	Hide the component's subcomponents		Volume
	Storage system		

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.

RAID6 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 the *AssuredSAN 4000 Series 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 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.

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:

- EC log, which includes EC debug data, EC revisions, and PHY statistics
- SC debug log and controller event log
- SC crash logs, which include the SC boot log
- 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 on page 61. For information about using the FTP interface, see Using FTP to download logs and update firmware on page 101.

Related topics

- Configuring email notification on page 32 (for push mode)
- Configuring SNMP notification on page 32 (for pull mode)
- Changing management interface settings on page 31 (to enable SNMP or SMI-S for pull mode)
- Enabling/disabling managed logs on page 43
- Testing notifications on page 65

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 (B)
- Data Throughput (B/s)
- Average Response Time (µs)

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 *AssuredSAN 4000 Series CLI Reference Guide*.

Related topics

- Vdisk performance on page 77 (to view historical performance statistics for a vdisk)
- Disk performance on page 83 (to view historical performance statistics for a disk)
- Resetting historical disk-performance statistics on page 65
- Saving historical disk-performance statistics on page 65

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, displaying Step 1 of 8: Introduction.
2. Review the text, then click **Next** to continue.

Step 2: Changing default passwords

In the wizard panel, Step 2 of 8: Password Setup displays.

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

In the wizard panel, Step 3 of 8: Network Configuration displays.

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:

- 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 by selecting DHCP in the IP address source list, 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 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 IP address source to **manual**.
3. Set the values for each controller. You must set a unique IP address for each network port. All fields are required.
4. Click **Next** to continue.

Step 4: Enabling system-management services

In the wizard panel, Step 4 of 8: Enable System-Management Services displays.

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

- WBI. The primary interface for managing the system. You can enable use of HTTP, or HTTPS for increased security, or both.
- CLI. An advanced user interface for managing the system. You can enable use of Telnet, or SSH for increased security, or both.
- SMI-S. Used for management of the system through your network. You can enable use of secure (encrypted) or unsecure (unencrypted) SMI-S:
  - Encrypted. Select this checkbox to enable encryption. 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.
  - Unencrypted. Clear this checkbox to disable encryption. 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.

  The SMI-S is a 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.

- Other Interfaces. Other methods of interaction with the system. You can enable FTP, SNMP, Service Debug, In-band SES Capability, or all.
  - FTP. A secondary interface for installing firmware updates and downloading logs.
  - SNMP. Used for monitoring of the system through your network.
  - Service Debug. Used for technical support only.
  - In-band SES Capability. Used for in-band monitoring of system status based on SES data. In-band management interfaces operate through the data path and can slightly reduce I/O performance.

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 on page 13.

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

In the wizard panel, Step 5 of 8: System Information displays.

1. Enter a **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.
2. Click **Next** to continue.

Step 6: Configuring event notification

In the wizard panel, Step 6 of 8: Configure Event Notification displays.

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

1. In the Email 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), which disables email notification.
   - **SMTP Server address** (required). 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 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** (required). 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`.

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), 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 letters, numbers, hyphens, and underscores; 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 letters, numbers, hyphens, and underscores; 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

In the wizard panel, Step 7 of 8: Port Configuration displays.

To enable the system to communicate with hosts, you must configure the system's host-interface options. There are options for FC ports. For SAS ports there are no host-interface options; click **Next** to continue.

For FC ports you can set these options:

- Speed. Can be set to auto, which auto-negotiates the proper link speed with the host, or to 2 Gbit (Gbit/s), 4 Gbit, or 8 Gbit. 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 for testing, or you need to specify a mutually supported speed for more than two FC devices connected in an arbitrated loop.
- Connection Mode. Can be set to loop, point-to-point, or auto. Loop protocol can be used in a physical loop or in a direct physical connection between two devices. Point-to-point protocol can only be used on a direct physical connection between exactly two devices. Auto sets the mode based on the detected connection type.
- Loop IDs can be set, per controller, to use soft or hard target addressing:
  - Soft target addressing enables a LIP to determine the loop ID. Use this setting if the loop ID is permitted to change after a LIP or power cycle.
  - Hard target addressing requests a specific loop ID that should remain after a LIP or power cycle. If the port cannot acquire the specified ID, it is assigned a soft target address. Use this option if you want ports to have specific addresses, if your system checks addresses in reverse order (lowest address first), or if an application requires that specific IDs be assigned to recognize the controller.

To change FC host-interface settings

1. For controller host ports that are attached to hosts:
   - Set the speed to the proper value to communicate with the host.
   - Set the connection mode.
2. For each controller, set the loop ID to use soft or hard target addressing. To use soft target addressing, select (enable) **Soft?**. To use hard target addressing, clear (disable) **Soft?** and enter an address in the

range 0–125. You cannot set the same hard target address for both controllers. An asterisk indicates that the value is required.

3. Click **Next** to continue.

Step 8: Confirming configuration changes

In the wizard panel, Step 8 of 8: Confirm the Configuration Changes displays with a table listing all the values selected in the wizard.

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.

NOTE: If you changed a controller's FC loop ID setting, you must restart the SC for the change to take effect.

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:

- WBI. The primary interface for managing the system. You can enable use of HTTP, of HTTPS for increased security, or both.
- CLI. An advanced user interface for managing the system. You can enable use of Telnet, of SSH for increased security, or both.
- SMI-S. Used for management of the system through your network. You can enable use of secure (encrypted) or unsecure (unencrypted) SMI-S:
  - Encrypted. Select this checkbox to enable encryption. 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.
  - Unencrypted. Clear this checkbox to disable encryption. 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.
- Other Interfaces. Other methods of interaction with the system. You can enable FTP, SNMP, Service Debug, In-band SES Capability, or all.
  - FTP. A secondary interface for installing firmware updates and downloading logs.
  - SNMP. Used for monitoring of the system through your network.
  - Service Debug. Used for technical support only.
  - In-band SES Capability. Used for in-band monitoring of system status based on SES data. In-band management interfaces operate through the data path and can slightly reduce I/O performance.

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 on page 13.

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 click 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 on page 24.

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), which disables email notification.
   - **SMTP Server address** (required). 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** (required). 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 65.

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 selected (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.
3. Click **Apply**.
4. Enable log management as described on page 43.
5. Send a test message to the configured destination as described on page 65.

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), 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 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 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 65.

Configuring syslog notification

To configure syslog notification of events

1. In the Configuration View panel, right-click the system and select **Configuration > Services > Syslog 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). None disables syslog notification.
   - **Syslog Server IP Address**. IP address of the syslog host system.
   - **Syslog Server Port Number**. Number port of the syslog host system.
3. Click **Apply**.
4. Optionally, send a test message to the configured destinations as described on page 65.

Configuring user accounts

Adding users

You can create either a general user that can access the WBI, 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** (required). 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.

NOTE: The user name `admin` is reserved for internal use.

   - **Password** (required). A password is case sensitive and can include a maximum of 32 bytes using characters except a backslash, comma, or double quote.
   - SNMPv3 User. Select **Standard User**.
   - **User Roles**. Select Monitor or Manage. You cannot change the roles of user `manage`.
     - **Monitor**. A user with role Monitor may view system settings but cannot change them.
     - **Manage**. A user with role Manage may view and change system settings.
   - **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**. The unit for display of storage-space sizes: **Auto**, **TB**, **GB**, or **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**. Specifies the scale to use for temperature values: **Celsius** or **Fahrenheit**.
- **Auto Sign Out**. Select 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**.

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** (required). 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.

   NOTE: The user name `admin` is reserved for internal use.

   - **Password** (required). A password is case sensitive and can include a maximum of 32 bytes using characters except a backslash, comma, 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.
   - SNMPv3 User. 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**). Authentication uses the user password.
   - **Privacy Type**. Select whether to use **DES** or **AES** encryption, or no encryption (**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, 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, 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, or double quote.
   - **User Roles**. Select Monitor or Manage. You cannot change the roles of user `manage`.
     - **Monitor**. A user with role Monitor may view system settings but cannot change them.
     - **Manage**. A user with role Manage may view and change system settings.
   - **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**. The unit for display of storage-space sizes: **Auto**, **TB**, **GB**, or **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**. Specifies the scale to use for temperature values: **Celsius** or **Fahrenheit**.
   - **Auto Sign Out**. Select 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**.
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, 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**). Authentication uses the user password.
- **Privacy Type**. Select whether to use **DES** or **AES** encryption, or no encryption (**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, 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**. 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 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 click 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 on page 22.

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**.
   - **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, and optionally minutes, from UTC. For example: the Pacific Time Zone offset is -8 during Pacific Standard Time or -7 during Pacific Daylight Time; the offset for Bangalore, India is +5:30.

- **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, you must configure the system's host-interface options. There are options for FC ports but not for SAS ports.

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 **Speed** to **auto**, which auto-negotiates the proper link speed with the host, or to **2 Gb** (Gbit/s), **4 Gb**, or **8 Gb**. 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 for testing, or if you need to specify a mutually supported speed for more than two FC devices connected in an arbitrated loop.
3. Set **Connection Mode** to **loop**, **point-to-point**, or **auto**. Loop protocol can be used in a physical loop or in a direct physical connection between two devices. Point-to-point protocol can only be used on a direct physical connection between exactly two devices. Auto sets the mode based on the detected connection type.
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:

- 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 try to 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 **IP address**, **IP mask**, and **Gateway** for each controller. You must set a unique IP address for each network port. All fields are required.
5. Record the IP values you assign.
6. Click **Apply**.
7. Sign out and try to 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

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 **SMART Configuration** 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.
   - **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 DSD for available disks and global spares

The 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 DSD for all disks on page 39. To configure DSD for a vdisk, see Configuring DSD for a vdisk on page 44.

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 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)**, which is the period of inactivity after which available disks and global spares automatically spin down, from 1–360 minutes.
3. Click **Apply**. When processing is complete a success dialog appears.
4. Click **OK**.

Scheduling DSD for all disks

For all disks that are configured to use 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 DSD for a vdisk on page 44. To configure DSD for available disks and global spares, see Configuring DSD for available disks and global spares on page 39.

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 system and select **Configuration > Advanced Settings > Disk**.
2. Set the options:
   - Select the **Drive Spin Down Suspend Period** option.
   - Set the **Time to Suspend** and the **Time to Resume** DSD 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 the **Exclude Weekend Days from Suspend Period** option.

3. Click **Apply**. When processing is complete a success dialog appears.
4. Click **OK**.

Configuring the EMP polling rate

You can change the interval at which the storage system polls each attached EMP for status changes.

- 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.
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. 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 **Sync Cache Mode** to either:
   - **Immediate**. Good status is returned immediately and cache content is unchanged.
   - **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 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 **Missing LUN** Response 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.
   - **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.

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' Independent Cache Performance Mode setting

When controller failover is enabled, data in a controller's write-back cache is mirrored to the partner controller. You can enable Independent Cache Performance Mode, in which controller failover is disabled and data in a controller's write-back cache is not mirrored to the partner controller. This improves write performance at the risk of losing unwritten data if a controller failure occurs while there is data in controller cache.

To change Independent Cache Performance Mode settings

1. In the Configuration View panel, right-click the system and select **Configuration > Advanced Settings > Cache**.
2. Either select (enable) or clear (disable) the I**ndependent Cache Performance Mode** option.
3. Click **Apply**.

Changing AWT 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 on page 18. You can also specify actions for the system to take when write-through caching is triggered.

To change AWT 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.
   - **Cache Power**. Changes to write-through if cache backup power is not fully charged or fails.
   - **CompactFlash**. Changes to write-through if CompactFlash memory is not detected during POST, fails during POST, or fails while the controller is under operation.
   - **Power Supply Failure**. Changes to write-through if a PSU fails.
   - **Fan Failure**. Changes to write-through if a cooling fan fails.
   - **Overtemperature Failure**. Forces a controller shutdown if a temperature is detected that exceeds system threshold limits.
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.
   - **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.
4. Click **Apply**.

Configuring PFU

In a system in which PFU is enabled, when you update firmware on one controller, the system automatically updates the partner controller.

IMPORTANT: Disable PFU only if requested by a service technician.

To change the PFU 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 for SAS disks.

TIP: If you choose to disable background vdisk scrub, you can still scrub a selected vdisk by using **Tools > Media Scrub Vdisk** (page 67).

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.
   - Set the **Vdisk Scrub Interval (hours)**, which is the interval between background vdisk scrub finishing and starting again, from 1–360 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. The interval between background disk scrub finishing and starting again is 24 hours.

Enabling background disk scrub is recommended for SAS disks.

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.
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 **Utility Priority** 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.
   - **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 on page 24. Before enabling log management, configure log destinations as described in Configuring email notification on page 32 and Configuring SNMP notification on page 32. After enabling log management, you can test it as described in Testing notifications on page 65.

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.
3. Click **Apply**.

Configuring an existing 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.

If a disk in the vdisk fails, a dedicated spare is automatically used to reconstruct the vdisk. A fault-tolerant vdisk other than RAID6 becomes Critical when one disk fails. A RAID6 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 RAID50 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 Tabular enclosure list or Graphical enclosure 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 20 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.

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 click Yes, a processing dialog appears. When processing is complete a success dialog appears.
5. Click **OK**.

Configuring DSD for a vdisk

The 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 DSD for all disks on page 39. To configure DSD for available disks and global spares, see Configuring DSD for available disks and global spares on page 39.

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)**, 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 an existing 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

△ **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**.
   - **Write Optimization**. Select **Standard**, **No-mirror**, or **Atomic Write**.
   - **Read Ahead Size**. Select **Default**, **Disabled**, **Maximum**, or a specific size (64, 128, 256, or 512 KB; 1, 2, 4, 8, 16, or 32 MB).
3. Click **Modify Cache Settings**.

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, displaying Step 1 of 6: Introduction.
2. Review the text, then click **Next** to continue.

Step 2: Specifying the vdisk name and RAID level

In the wizard panel, Step 2 of 6: Name and RAID Level displays.

For information on vdisks, see About vdisks on page 14.

To create a vdisk

1. Set the options:
   - **Vdisk name** (required). This field is populated with a default name. You may 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 20 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. **Auto** automatically assigns the owner to load-balance vdisks between controllers.
   - **RAID level** (required). Select a RAID level for the vdisk.
   - **Number of sub-vdisks**. This field is populated with a default number for a RAID10 or RAID50 vdisk. You may optionally change the number of sub-vdisks that the vdisk should contain.
   - **Chunk size**. For RAID 3, 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 RAID50, this option sets the chunk size of each RAID5 sub-vdisk. The chunk size of the RAID50 vdisk is calculated as: *configured-chunk-size* x (*subvdisk-members* - 1). For NRAID and RAID1, chunk size has no meaning and is therefore disabled.
2. Click **Next** to continue.

Step 3: Selecting disks

In the wizard panel, Step 3 of 6: Select Disks displays.

Select disks to include in the vdisk. The Disk Selection Sets table has one row for each sub-vdisk in a RAID10 or RAID50 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 on page 22.

The Tabular Enclosures Front View 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 all 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

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

In the wizard panel, Step 4 of 6: Define Volumes displays.

For information on volumes, see About volumes on page 16.

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** (required). 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 is the total space 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

In the wizard panel, Step 5 of 6: Set the Base Mapping for Volumes displays.

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

- A 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 Volume1, 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: **no-access** (masked), **read-only**, or **read-write**.
5. Click **Next** to continue.

Step 6: Confirming vdisk settings

In the wizard panel, Step 6 of 6: Confirm the Vdisk Settings displays with a table listing all the values selected in the wizard.

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** (required). This field is populated with a default name. You may 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 20 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. Auto automatically assigns the owner to load-balance vdisks between controllers.
   - **RAID Level**. Select a RAID level for the vdisk.
   - **Number of Sub-vdisks**. For a RAID10 or RAID50 vdisk, optionally change the number of sub-vdisks that the vdisk should contain.
   - **Chunk size**. For RAID 3, 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 RAID50, this option sets the chunk size of each RAID5 sub-vdisk. The chunk size of the RAID50 vdisk is calculated as: *configured-chunk-size* x (*subvdisk-members* - 1). For NRAID and RAID1, chunk size has no meaning and is therefore disabled.
   - **Online Initialization**. If this option is enabled, 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 this option is disabled, 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-vdisk in a RAID10 or RAID50 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 on page 22.

   The Tabular Enclosures Front View 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 all disk information graphically. 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, right-click the system or **Vdisks** and then select **Provisioning > Delete Vdisks**.
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 click 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.

To delete a vdisk

1. Verify that hosts are not accessing volumes in the vdisk that you want to delete.
2. In the Configuration View panel, right-click a vdisk and then select **Provisioning > Delete Vdisks**.
3. Click **Delete Vdisk(s)**. A confirmation dialog appears.
4. Click **Delete** to continue; otherwise, click **Cancel**. If you click Delete, a processing dialog appears. If the task succeeds, an overview panel and a success dialog appear.
5. Click **OK**. As processing completes, the deleted items are removed from the Configuration View panel.

Managing global spares

You can designate a maximum of 64 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 RAID50 vdisks, if more than one sub-vdisk becomes critical, reconstruction and use of spares occur in the order sub-vdisks 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.
   - The Tabular Enclosures Front View tab shows all available disks in all enclosures in a table. The Graphical tab shows all disk information graphically. Disks you select are highlighted and color-coded to match the rows in the Disk Selection Sets table.

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.

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** (required). This field is populated with a default name. You may optionally change the base name for the volumes. The volume names will consist of the base name and a number that increments from 0000. 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 Volume0000, if Volume0001 already exists, the second volume is named Volume0002. 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:
     - **LUN** (required). 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.
     - **Access**. Select the access level that hosts will have to the volumes.
     - 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.

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** (required). This field is populated with a default name. You may optionally change the default 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.
   - **Size**. Optionally change the default size, which is all free space in the vdisk.
   - **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.

Deleting volumes

You can use the Delete Volumes panel to delete volumes.

△ **CAUTION:** Deleting a volume removes its mappings 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, right-click the system, **Vdisks,** a vdisk, or a volume and then select **Provisioning > Delete Volumes**. A table displays, showing all the volumes for the selected item.
3. In the table, 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 click 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. 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: 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**.
2. A table displays, showing all the volumes for the selected vdisk. 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. Select **Map**.

4. Either:
   - Map the volumes to all hosts by setting a **Starting LUN** (required), selecting ports, and setting **Access** to **read-only** or **read-write**.
   - Mask the volumes from all hosts by setting a **Starting LUN** (required), selecting ports, and setting **Access** to **no-access**. Setting the default mapping to **no-access** will result in the LUN mapping being removed.
5. Click **Apply**. A message specifies whether the change succeeded or failed.
6. 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: 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**.
2. A table displays, showing all the volumes for the selected vdisk. 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 Maps for Selected Volumes table, select the host to change access for.
4. Select **Map**.
5. 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**.
6. Click **Apply**. A message specifies whether the change succeeded or failed.
7. 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: 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** (required) 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: 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.
- 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** (required) 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** (required) 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**.
2. A table displays, showing all the volumes for the selected vdisk. 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 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.

To expand a volume

1. In the Configuration View panel, right-click a 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.

Adding a host

To add a host

1. Determine the host's WWN.
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)** (required). Enter the host's WWN.
   - **Host Name** (required). This field is populated with a default name. You may 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, right-click the system or **Hosts** and then 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 click 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.

To remove a host

1. Verify that the host you want to remove is not accessing volumes.
2. In the Configuration View panel, right-click a host and then select **Provisioning > Remove Host**.
3. Click **Remove Host(s)**. A confirmation dialog appears.
4. Click **Remove** to continue; otherwise, click **Cancel**. If you click Remove, a processing dialog appears. If the task succeeds, an overview panel and a success dialog appear.
5. Click **OK**. As processing completes, the deleted item is 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 on page 53.

△ **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: 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.

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.

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 *AssuredSAN 4000 Series Service Guide*, and Removing a vdisk from quarantine on page 68.
- 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 *AssuredSAN 4000 Series Service Guide* and information about the `clear cache` command in the *AssuredSAN 4000 Series 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 71).

Updating controller-module firmware

A controller enclosure can contain one or two controller modules. Both controllers should run the same firmware version. You can update the firmware in each controller module by loading a firmware file obtained from the enclosure vendor.

If the 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 *AssuredSAN 4000 Series Release Notes*.

To update controller-module firmware

1. Obtain the appropriate firmware file and download it to your computer or network.
2. Restart the 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 on page 63.
3. 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.
4. Click **Browse** and select the firmware file to install.

5. 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 EMP to be updated. This typically takes 3 minutes for each EMP in a drive enclosure.

   If the SC 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 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.

6. 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 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. 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.
3. Select the expansion modules to update.
4. Click **Browse** and select the firmware file to install.
5. Click **Install Expansion-Module Firmware File**. A dialog box shows firmware-update progress.

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

It typically takes 3 minutes to update each EMP in a drive enclosure. Wait for a message that the code load has completed.

6. 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.

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**. A table titled Current Versions (Revisions) of All Disk Drives displays 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 on page 63.
   b. Power cycle all enclosures as described in the *AssuredSAN 4000 Series Setup Guide*.
9. Verify that each disk has the correct firmware revision.

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 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.
   b. Enter **Comments**, describing the problem and specifying the date and time when the problem occurred. This information helps service personnel when they analyze the log data. Comment text can be 500 bytes long.
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 an FC host port configured to use FC-AL topology, a reset issues a 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 to reset. For example, to reset controller A port 1, select **A1**.
3. Click **Reset Host Port**.

Rescanning disk channels

A rescan forces a rediscovery of disks and enclosures in the storage system. If both SCs are online and able to communicate with both expansion modules in each connected enclosure, rescan also reassigns the enclosure IDs of attached enclosures based on controller A's enclosure cabling order. A manual rescan may be needed after system power-up to display enclosures 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 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, in the CLI use the `restore defaults` command, as described in the *AssuredSAN 4000 Series 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 SC processor or the MC processor. A shut down affects both processors.

Restarting

If you restart an SC, 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 controller restarts. The MC is not restarted so it can provide status information to external interfaces.

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

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

NOTE: When an SC is restarted, live performance statistics that it recorded will be reset; historical performance statistics are not affected. 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 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 click Yes, a second confirmation dialog appears.
5. Click **Yes** to continue; otherwise, click **No**. If you click Yes, a message describes restart activity.

Shutting down

Shutting down the SC 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 SC 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 SCs are shut down, but information shown might be invalid.

To perform a shut down

1. In the Configuration View panel, right-click the system and select **Tools > Shut Down or Restart Controller**.
2. In the main panel, set the options:
   - **Operation**. Select **Shut down**.
   - **Controller Type**. Select **Storage**.
   - **Controller**. Select the controller to restart: **A**, **B**, or **Both**.
3. Click **Shut down now**. A confirmation dialog appears.

4. Click **Yes** to continue; otherwise, click **No**. If you click Yes, a second confirmation dialog appears.
5. Click **Yes** to continue; otherwise, click **No**. If you click 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 on page 24.

To test event notification

1. In the Configuration View panel, right-click the 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 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.

Checking system links

You can check the links between the partner controllers to see the host ports and their connections.

To check system links

1. In the Configuration View panel, right-click the system and select **Tools > Check Local System Link**.
2. Click **Check Links**.
   A table displays the host ports and their connections.

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 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 click 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 CSV 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 *AssuredSAN 4000 Series 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 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.

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 RAID1 vdisk.

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 Confirm Operation 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 RAID0.

TIP: Media Scrub Vdisk (page 67) operates similarly to Verify Vdisk but can find and fix media errors for any RAID level, including NRAID and RAID0.

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 on page 38).

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 on page 42) 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 on page 42.

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

Removing a vdisk from quarantine

CAUTION: If a vdisk is removed from quarantine and does not have enough disks to continue operation, its status will change to OFFL and its data cannot be recovered. To continue operation, a RAID3 or RAID5 vdisk can have only one inaccessible disk; a RAID6 vdisk can have only one or two inaccessible disks; a RAID10 or RAID50 vdisk can have only one inaccessible disk per sub-vdisk. For example, a 16-disk RAID10 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 RAID0 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 RAID6 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 down disks): The vdisk is degraded or online with at least one inaccessible disk. The vdisk can be accessed and remains fault tolerant. For example, one disk is inaccessible in a RAID6 vdisk.
- QTCR (quarantined critical): The vdisk is critical with at least one inaccessible disk. The vdisk can be accessed. For example, two disks are inaccessible in a RAID6 vdisk or one disk is inaccessible for other fault-tolerant RAID levels.
- QTOF (quarantined offline): The vdisk is fault tolerant with multiple inaccessible disks causing user data to be incomplete, or is an NRAID or RAID0 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. 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 down disks). 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 (RAID6 only) and its status changes to FTOL, CRIT, or FTDN. For status descriptions, see Vdisk properties on page 77.

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.
- Count.
- Capacity.
- Storage Space.

The table gives the following information about the individual components:

- The system's capacity and total storage space.
- The health, quantity, capacity, and space usage of enclosures, disks, vdisks, and volumes.
- Configuration limits and versions of controller firmware and hardware.

For descriptions of storage-space color codes, see About storage-space color codes on page 22.

Select a component to see more information about it.

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.
- System Contact.
- System Location.
- System Information. 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: Dual-port SAS.
  - sSAS: Dual-port SAS SSD.
  - SAS MDL: SAS Midline.
- 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.

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.
- 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; however, the vdisk could be accessed. For instance, one disk is missing from a mirror or RAID5.
    - QTDN: Quarantined with down disks. The vdisk is offline and quarantined because at least one disk is missing; however, the vdisk can be accessed and is fault tolerant. For instance, one disk is missing from a RAID6.
    - 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. Disk type values are described in Type on page 72.

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.

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 schedule name and schedule specifications:

- Schedule Name.
- Schedule Specifications. The start date 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.

If there are no schedules in the system, no tables display.

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 the system supports.

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 *AssuredSAN 4000 Series Service 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 *AssuredSAN 4000 Series Service 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 displays the following information:

- 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 on page 22.

The Vdisks table displays 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. Status values are described in Status on page 73.
- Disk Type. Disk type values are described in Type on page 72.
- 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 Vdisk Overview table shows:

- Health.
  - OK
  - Degraded
  - Fault
  - Unknown
- Component. Vdisk, disks, volumes.
- Count.
- Capacity.
- Storage space. For descriptions of storage-space color codes, see About storage-space color codes on page 22.

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 select the **Properties** tab, the Properties for *Vdisk* table shows:

- Health.
  - OK
  - Degraded
  - Fault
  - Unknown
- Health Reason. If a vdisk's health is not OK, this entry lists the reasons for the health state. If a vdisk's health is OK, no information is displayed.
- Health Recommendation. If a vdisk's health is not OK, this entry lists recommendations for correcting the health. If a vdisk's health is OK, no information is displayed.
- 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.
- 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. The amount of contiguous data that is written to a vdisk member before moving to the next member of the vdisk.
- Created. Date and time the vdisk was created.
- Minimum Disk Size.
- Status. Status values are described in Status on page 73
- Current Job.
  - 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. Enabled or disabled.

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".

Vdisk performance

When you select the Vdisk component and select 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 on page 83) 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 GB.
- 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 MB/s.
- 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 milliseconds. 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 81). To view live (non-historical) performance statistics for one more vdisks, in the CLI use the `show vdisk-statistics` command.

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 RAID10 or RAID50, the sub-vdisk that the disk is in; for other RAID levels, the disk's RAID level; or SPARE.
- Disk Type. Disk type values are described in Type on page 72.
- 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. Shows whether the disk is healthy or has a problem.
   OK
   Degraded
   Fault
   N/A
   Unknown
  If the disk's health is not OK, view health details in the Enclosure Overview panel (page 81).
- Name. System-defined disk name using the format Disk-*enclosure-number.disk-slot-number*.
- Type. Disk type values are described in Type on page 72.
- 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. Current Job values are described in Current Job on page 73.
- 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. State values are described in State. Shows how the disk is used: on page 79

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.

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.
- 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 on page 22.

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 on page 18.
- Cache Optimization. Standard or No-mirror. See Optimizing read-ahead caching on page 19.
- Read Ahead Size. See Optimizing read-ahead caching on page 19.
- Type. Standard (typical) volume.
- 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.

Maps 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.
- Host Name. User-defined nickname for 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.

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.
- Host 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 host.
- Host Type.
  - If the host was discovered and its entry was automatically created, its host-interface type: FC or 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:

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

Maps 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.

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
  - N/A

- 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 EMP in controller module A's EC.
- EMP B Revision. Firmware revision of the EMP in controller module B's EC.
- 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 select the **Properties** tab, 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 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.
- Enclosure ID.
- Slot.
- How Used. How Used values are described in How Used on page 72.
- Type. Disk type values are described in Type on page 72.

- 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. Current Job values are described in Current Job on page 73.
- SMART. Shows whether SMART is enabled. For more information, see Configuring SMART on page 38.
- 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.

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 (B). Shows the amounts of data read and written and the combined total over the sampling time period. The base unit is bytes.

- Data Throughput (B/s). 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 (µs). Shows the average response times for reads and writes and the combined overage over the sampling time period. The base unit is microseconds.
- Average I/O Size (B). 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 76. 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 PSU, 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. The port is operating normally.
  - Degraded. The port's operation is 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: host port properties

When you select a 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.
- Ports.
- Media.
- Target ID.
- Configured Speed.
- Actual Speed.

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.

A SNMP reference

This appendix describes the SNMP capabilities that 4000 Series storage systems support. This includes standard MIB-II, the FA SNMP MIB version 2.2 objects, and enterprise traps.

4000 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 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 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

4000 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 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 OID (`sysObjectID`) is based on the vendor name followed by ".2." and the identifier for the particular product model. For example, the OID for 4000 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 FA 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 vdisks, and the failure of any fault-tolerant 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), and a "sensor" in the sensor table.

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

The port section is only relevant to products with FC 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 4000 Series storage system. Unless specified otherwise, objects are *not* settable.

Table 12 FA MIB 2.2 objects, descriptions, and values

Object	Description	Value
`revisionNumber`	Revision number for this MIB	0220
`uNumber`	Number of connectivity units present	1
`systemURL`	Top-level URL of the device; for example, `http://10.1.2.3`. If a web server is not present on the device, this string is empty in accordance with the FA MIB2.2 Spec.	Default: `http://10.0.0.1`
`statusChangeTime`	`sysuptime` timestamp of the last status change event, in centiseconds. `sysuptime` starts at 0 when the SC boots and keeps track of the up time. `statusChangeTime` is updated each time an event occurs.	`0` at startup
`configurationChangeTime`	`sysuptime` timestamp of the last configuration change event, in centiseconds. `sysuptime` starts at 0 when the SC boots and keeps track of the up time. `configurationChangeTime` is updated each time an event occurs.	`0` at startup
`connUnitTableChangeTime`	`sysuptime` timestamp of the last update to the `connUnitTable` (an entry was either added or deleted), in centiseconds	`0` always (entries are not added to or deleted from the `connUnitTable`)

Table 12 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 ID 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
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
connUnitStatus ChangeTime	Same as statusChangeTime	Same as statusChangeTime
connUnit ConfigurationChange Time	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 12 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 on page 95
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.
connUnitRevs Description	Description of a component to which the revision corresponds	See External details for connUnitRevsTable on page 95
connUnitSensorTable	**Includes the following objects as specified by the FA MIB2.2 Spec**	
connUnitSensorUnit Id	connUnitId of the connectivity unit that contains this sensor table	Same as connUnitId
connUnitSensorIndex	Unique value for each connUnitSensorEntry between 1 and connUnitNumSensors	See External details for connUnitSensorTable on page 96
connUnitSensorName	Textual ID of the sensor intended primarily for operator use	See External details for connUnitSensorTable on page 96
connUnitSensor Status	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
connUnitSensor Message	Description the sensor status as a message	connUnitSensorName followed by the appropriate sensor reading. Temperatures display in both Celsius and Fahrenheit; for example, CPU Temperature (Controller Module A): 48C 118F). Reports “Not installed” or “Offline” if data is not available.
connUnitSensorType	Type of component being monitored by this sensor	See External details for connUnitSensorTable on page 96
connUnitSensor Characteristic	Characteristics being monitored by this sensor	See External details for connUnitSensorTable on page 96

Table 12 FA MIB 2.2 objects, descriptions, and values (continued)

Object	Description	Value
connUnitPortTable	**Includes the following objects as specified by the FA MIB2.2 Spec**	
connUnitPortUnitId	connUnitId of the connectivity unit that contains this port	Same as connUnitId
connUnitPortIndex	Unique value for each connUnitPortEntry between 1 and connUnitNumPorts	Unique value for each port, between 1 and the number of ports
connUnitPortType	Port type	not-present(3), or n-port(5) for point-to-point topology, or l-port(6)
connUnitPortFC ClassCap	Bit mask that specifies the classes of service capability of this port. If this is not applicable, returns all bits set to zero.	FC ports return 8 for class-three
connUnitPortFC ClassOp	Bit mask that specifies the classes of service that are currently operational. If this is not applicable, returns all bits set to zero.	FC ports return 8 for class-three
connUnitPortState	State of the port hardware	unknown(1), online(2), offline(3), bypassed(4)
connUnitPortStatus	Overall protocol status for the port	unknown(1), unused(2), ok(3), warning(4), failure(5), notparticipating(6), initializing(7), bypass(8)
connUnitPort TransmitterType	Technology of the port transceiver	unknown(1) for FC ports
connUnitPortModule Type	Module type of the port connector	unknown(1)
connUnitPortWwn	FC WWN of the port if applicable	WWN octet for the port, or empty string if the port is not present
connUnitPortFCId	Assigned FC ID of this port	FC ID of the port All bits set to 1 if the FC ID is not assigned or if the port is not present
connUnitPortSn	Serial number of the unit (for example, for a GBIC). If this is not applicable, returns an empty string.	Empty string
connUnitPort Revision	Port revision (for example, for a GBIC)	Empty string
connUnitPortVendor	Port vendor (for example, for a GBIC)	Empty string
connUnitPortSpeed	Speed of the port in KB/s (1 KByte = 1000 Byte)	Port speed in KB/s, or 0 if the port is not present
connUnitPortControl	Not supported	invalid(2) for an SNMP GET operation and not settable through an SNMP SET operation
connUnitPortName	String describing the addressed port	See External details for connUnitPortTable on page 98
connUnitPort PhysicalNumber	Port number represented on the hardware	Port number represented on the hardware
connUnitPortStat Object	Not supported	0 (No statistics available)

Table 12 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: DDMMYYYY HHMMSS	0 for logged events that occurred prior to or at startup
connUnitSEventTime	sysuptime timestamp when the event occurred	0 at startup
connUnitEvent Severity	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
connUnitPortStat FabricTable	Not supported	N/A
connUnitPortStat SCSITable	Not supported	N/A
connUnitPortStatLAN Table	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 12 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, DH4000
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 12.

External details for connUnitRevsTable

Table 13 connUnitRevsTable index and description values

connUnitRevsIndex	connUnitRevsDescription
1	SC processor Type (I/O Manager-A)
2	Bundle Version (I/O Manager-A)
3	Build Date (I/O Manager-A)
4	SC Code Version (I/O Manager-A)
5	SC Code Baselevel (I/O Manager-A)
6	Memory Controller FPGA Code Version (I/O Manager-A)
7	SC Loader Code Version (I/O Manager-A)
8	CAPI Version (I/O Manager-A)
9	MC Code Version (I/O Manager-A)
10	MC Loader Code Version (I/O Manager-A)
11	EC Code Version (I/O Manager-A)
12	CPLD Code Version (I/O Manager-A)

Table 13 connUnitRevsTable index and description values (continued)

connUnitRevsIndex	connUnitRevsDescription
13	Hardware Version (I/O Manager-A)
14	Host Interface Module Version (I/O Manager-A)
15	Host Interface Module Model (I/O Manager-A)
16	Backplane Type (I/O Manager-A)
17	Host Interface Hardware (Chip) Version (I/O Manager-A)
18	Disk Interface Hardware (Chip) Version (I/O Manager-A)
19	SC processor Type (I/O Manager-A)
20	Bundle Version (I/O Manager-B)
21	Build Date (I/O Manager-B)
22	SC Code Version (I/O Manager-B)
23	SC Code Baselevel (I/O Manager-B)
24	Memory Controller FPGA Code Version (I/O Manager-B)
25	SC Loader Code Version (I/O Manager-B)
26	CAPI Version (I/O Manager-B)
27	MC Code Version (I/O Manager-B)
28	MC Loader Code Version (I/O Manager-B)
29	EC Code Version (I/O Manager-B)
30	CPLD Code Version (I/O Manager-B)
31	Hardware Version (I/O Manager-B)
32	Host Interface Module Version (I/O Manager-B)
33	Host Interface Module Model (I/O Manager-B)
34	Backplane Type (I/O Manager-B)
35	Host Interface Hardware (Chip) Version (I/O Manager-B)
36	Disk Interface Hardware (Chip) Version (I/O Manager-B)

External details for connUnitSensorTable

Table 14 connUnitSensorTable index, name, type, and characteristic values

connUnitSensorIndex	connUnitSensorName	connUnitSensorType	connUnitSensor Characteristic
1	On-Board Temperature 1 Controller A	board(8)	temperature
2	On-Board Temperature 1 Controller B	board(8)	temperature
3	On-Board Temperature 2 Controller A	board(8)	temperature
4	On-Board Temperature 2 Controller B	board(8)	temperature
5	On-Board Temperature 3 Controller A	board(8)	temperature
6	On-Board Temperature 3 Controller B	board(8)	temperature
7	Disk Controller Temp Controller A	board(8)	temperature
8	Disk Controller Temp Controller B	board(8)	temperature

Table 14 connUnitSensorTable index, name, type, and characteristic values (continued)

connUnitSensorIndex	connUnitSensorName	connUnitSensorType	connUnitSensor Characteristic
9	Memory Controller Temp Controller A	board(8)	temperature
10	Memory Controller Temp Controller B	board(8)	temperature
11	Capacitor Pack Voltage Controller A	board(8)	power
12	Capacitor Pack Voltage Controller B	board(8)	power
13	Capacitor Cell 1 Voltage Controller A	board(8)	power
14	Capacitor Cell 1 Voltage Controller B	board(8)	power
15	Capacitor Cell 2 Voltage Controller A	board(8)	power
16	Capacitor Cell 2 Voltage Controller B	board(8)	power
17	Capacitor Cell 3 Voltage Controller A	board(8)	power
18	Capacitor Cell 3 Voltage Controller B	board(8)	power
19	Capacitor Cell 4 Voltage Controller A	board(8)	power
20	Capacitor Cell 4 Voltage Controller B	board(8)	power
21	Capacitor Charge Controller A	board(8)	other
22	Capacitor Charge Controller B	board(8)	other
23	Overall Unit Status: OK	enclosure(7)	other
24	Temperature Loc: upper IOM A	enclosure(7)	temperature
25	Temperature Loc: lower IOM B	enclosure(7)	temperature
26	Temperature Loc: left PSU	power-supply(5)	temperature
27	Temperature Loc: right PSU	power-supply(5)	temperature
28	Voltage 12V Loc: upper-IOM A	enclosure(7)	power
29	Voltage 5V Loc: upper-IOM A	enclosure(7)	power
30	Voltage 12V Loc: lower-IOM B	enclosure(7)	power
31	Voltage 5V Loc: lower-IOM B	enclosure(7)	power
32	Voltage 12V Loc: left-PSU	power-supply(5)	power
33	Voltage 5V Loc: left-PSU	power-supply(5)	power
34	Voltage 3.3V Loc: left-PSU	power-supply(5)	power
35	Voltage 12V Loc: right PSU	power-supply(5)	power
36	Voltage 5V Loc: right PSU	power-supply(5)	power
37	Voltage 3.3V Loc: right PSU	power-supply(5)	power
38	Current 12V Loc: upper-IOM A	enclosure(7)	current
39	Current 12V Loc: lower-IOM B	enclosure(7)	current
40	Current 12V Loc: left-PSU	power-supply(5)	current
41	Current 5V Loc: left-PSU	power-supply(5)	current
42	Current 12V Loc: right-PSU	power-supply(5)	current
43	Current 5V Loc: right-PSU	power-supply(5)	current

External details for connUnitPortTable

Table 15 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 on page 31.
2. Configure and enable SNMP traps; see Configuring SNMP notification on page 32.
3. Optionally, configure a user account to receive SNMP traps; see Configuring user accounts on page 33.

SNMP management

You can manage storage devices using SNMP with a network management system such as HP OpenView, HP SIM, or HP 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 on page 31. 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 on page 33. 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 4000 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 to download logs and update firmware

Although RAIDar is the preferred interface for downloading log data and historical disk-performance statistics and updating firmware, 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.

NOTE: Collecting log information has a negative impact on performance while collection is in progress.

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:

- User configuration settings from both controllers
- Event logs from both controllers
- SC logs from both controllers
- SC crash dumps from both controllers
- CAPI trace from the controller receiving the command
- MC log from the controller receiving the command
- Controller environment (including data about attached disks, enclosures, and so forth)

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 37.
   b. Verify that the system's FTP service is enabled; see Changing management interface settings on page 31.
   c. Verify that the user you will log in as has permission to use the FTP interface; see Modifying users on page 35.
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 on page 24.

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 37.
   b. Verify that the system's FTP service is enabled; see Changing management interface settings on page 31.
   c. Verify that the user you will log in as has permission to use the FTP interface; see Modifying users on page 35.
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 SC's four crash logs.
     - `ecdebug`: EC log.
     - `mc`: MC log.
     - `scdebug`: SC 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 CSV 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 *AssuredSAN 4000 Series 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 37.
   b. Verify that the system's FTP service is enabled; see Changing management interface settings on page 31.
   c. Verify that the user you will log in as has permission to use the FTP interface; see Modifying users on page 35.
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 *AssuredSAN 4000 Series Service Guide*, and Removing a vdisk from quarantine on page 68.
- If any unwritten cache data is present, firmware update will not proceed. Before you can update firmware, that data must be removed from cache. See information about event 44 in the *AssuredSAN 4000 Series Service Guide* and information about the `clear cache` command in the *AssuredSAN 4000 Series 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 71).

Updating controller-module firmware

A controller enclosure can contain one or two controller modules. Both controllers should run the same firmware version. You can update the firmware in each controller module by loading a firmware file obtained from the enclosure vendor.

If the 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 *AssuredSAN 4000 Series 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. Restart the 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 on page 63.
4. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the directory containing the firmware file to load.
5. Enter:

   ftp *controller-network-address*

   For example:

   ftp 10.1.0.9
6. Log in as an FTP user.
7. Enter:

   put *firmware-file* **flash**

   For example:

   put CF100R01-01.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 having downlevel CPLD firmware. If the controller enclosure has attached drive enclosures, allow additional time for each EMP to be updated. This typically takes 3 minutes for an EMP in each 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 SC 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.
8. Quit the FTP session.
9. 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 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. Open a Command Prompt (Windows) or a terminal window (UNIX) and navigate to the directory containing the firmware file to load.
5. Enter:

   ftp *controller-network-address*

   For example:

   ftp 10.1.0.9
6. Log in as an FTP user.
7. 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 1.5 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.

8. If you are updating specific expansion modules, repeat step 7 for each remaining expansion module that needs to be updated.
9. Quit the FTP session.
10. 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.

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, determine the enclosure number and slot number of the disk.
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 the *AssuredSAN 4000 Series Setup Guide*.
8. Verify that each disk has the correct firmware revision.

C Using SMI-S

This appendix provides information for network administrators who are managing the 4000 Series from a storage management application through the SMI-S. SMI-S is a 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:

- 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
- 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.
- 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
- 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
- Job Control subprofile
- Storage Enclosure subprofile (if expansion enclosures are attached)
- Disk Sparing 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 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, Dot Hill AssuredSAN 4000 Series storage systems), and provides the CIMOM with information about them. In theory, providers for multiple types of devices (for example, 4000 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 MC architecture, use limited disk space, use limited memory resources and be as fast as a proxy provider running on a server. The provider is an MC application and works by making MC CLI requests. An SMI-S cache caches these requests for 30 to 60 seconds. The disk space used is about 3 MB without qualifiers and 8 MB with qualifiers. 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 provider and uses this interface. To reduce the memory footprint, a third-party package called CIMPLE (www.simplewbem.org) is used. For more information on SFCB go to sblim.cvs.sourceforge.net/sblim/sfcb/README?view=markup.

About the 4000 Series SMI-S provider

The CF100 release passes all SMI-S 1.5 tests and is CTP 1.5 certified. Full provisioning is supported.

The 4000 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.

Each 4000 Series model is 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:

- 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. HP SIM determines which profiles it intends to manage, and then uses the CIM model to discover the actual configurations and capabilities.

Table 16 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 FC specific aspects of a target storage system.
SAS Target Ports subprofile	Models the SAS specific aspects of a target storage system.
Access Points subprofile	Provides addresses of 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.

Table 16 Supported SMI-S profiles (continued)

Profile/subprofile/package	Description
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.
Job Control subprofile	Provides the ability to monitor provisioning operations, such as creating volumes, 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.

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.

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 allow 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 on page 114.

The event categories in Table 17 pertain to FRU assemblies and certain FRU components.

Table 17 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

Table 17 CIM Alert indication events (continued)

FRU/Event category	Corresponding SMI-S class	Operational status values that would trigger alert conditions
Battery/Super Cap	DHS_SuperCap	Unknown, Error, OK
FC Port	DHS_FCPort	Stopped, OK
SAS Port	DHS_SASTargetPort	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. CQL and WQL are both supported, with some limitations to the CQL indication filter.

Table 18 Life cycle indications

Profile or subprofile	Element description and name	WQL or CQL
Block Services	SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_StoragePool Send life cycle indication when a vdisk is created or deleted.	Both
Block Services	SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_StorageVolume Send life cycle indication when a volume is created or deleted.	Both
Block Services	SELECT * FROM CIM_InstModification WHERE SourceInstance ISA CIM_LogicalDevice Send life cycle indication when disk drive (or any logical device) status changes.	Both
Disk Drive Lite	SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_DiskDrive Send life cycle indication when a disk drive is inserted or removed.	Both
Fan	SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_Fan Send life cycle indication when a fan is powered on or off.	Both
Job Control	SELECT * FROM CIM_InstModification WHERE SourceInstance ISA CIM_ConcreteJob AND SourceInstance.OperationalStatus=17 AND SourceInstance.OperationalStatus=2 Send life cycle indication when a create or delete operation completes for a volume or LUN.	WQL
Masking and Mapping	SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_AuthorizedSubject Send life cycle indication when a host privilege is created or deleted.	Both
Masking and Mapping	SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_ProtocolController Send life cycle indication when create/delete storage hardware ID (add/remove hosts).	Both
Masking and Mapping	SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_ProtocolControllerForUnit Send life cycle indication when a LUN is created, deleted, or modified.	Both
Multiple Computer System	SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_ComputerSystem Send life cycle indication when a controller is powered on or off.	Both
Multiple Computer System	SELECT * FROM CIM_InstModification WHERE SourceInstance ISA CIM_ComputerSystem AND SourceInstance.OperationalStatus <> PreviousInstance.OperationalStatus Send life cycle indication when a logical component degrades or upgrades the system.	WQL
Multiple Computer System	SELECT * FROM CIM_InstModification WHERE SourceInstance ISA CIM_RedundancySet AND SourceInstance.RedundancyStatus <> PreviousInstance.RedundancyStatus Send life cycle indication when the controller active-active configuration changes.	WQL

Table 18 Life cycle indications (continued)

Profile or subprofile	Element description and name	WQL or CQL
Target Ports	SELECT * FROM CIM_InstCreation WHERE SourceInstance ISA CIM_FCPort Send life cycle indication when a target port is created or deleted.	Both
Target Ports	SELECT * FROM CIM_InstModification WHERE SourceInstance ISA CIM_FCPort AND SourceInstance.OperationalStatus <> PreviousInstance.OperationalStatus Send life cycle indication when the status of a target port changes.	WQL

SMI-S configuration

In the default SMI-S configuration:

- The secure SMI-S protocol is enabled, which is the recommended protocol for SMI-S.
- The SMI-S interface is enabled for the `manage` user.

Table 19 lists the CLI commands relevant to the SMI-S protocol.

Table 19 CLI commands for SMI-S

Action	CLI command
Enable secure SMI-S port 5989 (and disable port 5988)	`set protocols smis enabled`
Disable secure SMI-S port 5989	`set protocols smis disabled`
Enable unsecure SMI-S port 5988 (and disable port 5989)	`set protocols usmis disabled`
Disable unsecure SMI-S port 5988 (and enable port 5989)	`set protocols usmis enabled`
Enable unsecure SMI-S port 5988	`set protocol usmis enabled`
See the current status	`show protocols`

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 on page 24.

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 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 1 step. The CreateStorageHardwareID and DeleteStorageHardwareID methods all 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 `manage` and type: `set protocol smis enabled`.
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 `smis` interface and set the user preferences to support the `smis` interface, if necessary. See SMI-S configuration on page 114 for instructions on how to add users. Also verify the supplied credentials.
Want to connect securely as user name `my_xxxx`.	Need to add user	Log in to the array as `manage`. Type `create user level manage my_xxxuser` and then type `set user my_xxxuser interfaces wbi,cli,smis`
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 wbemcli on a Linux system by typing `apt-get install wbemcli`. Type `wbemcli -nl -t -noverify ein 'https://manage:!manage@:5989/root/dhs:cim_computersystem'`

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 on page 24.

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:

- EC log, which includes EC debug data, EC revisions, and PHY statistics
- SC debug log and controller event log
- SC crash logs, which include the SC boot log
- 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 on page 102.

Log-file details

- SC debug-log records contain date/time stamps of the form *mm*/*dd hh*:*mm*:*ss*.
- SC crash logs (diagnostic dumps) are produced if the firmware fails. Upon restart, such logs are available, and the restart boot log is also included. The four most recent crash logs are retained in the storage system.
- When EC debug logs are obtained, EC revision data and SAS PHY statistics are also provided.
- MC debug logs transferred by the managed logs feature are for five internal components: `appsv`, `mccli`, `logc`, `web`, and `snmpd`. The contained files are log-file segments for these internal components and are numbered sequentially.

Storing log files

It is recommended to store log files hierarchically by storage-system name, log-file type, and date/time. Then, if historical analysis is required, the appropriate log-file segments can easily be located and can be linked together 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:

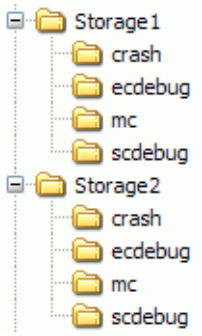

In push mode, when the administrator receives an email with an attached `ecdebug` file from Storage1, the administrator would open the attachment and unzip it into the ecdebug subdirectory of the `Storage1` directory.

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.
atomic write	A mode that guarantees if a failure (such as I/O being aborted or a controller failure) interrupts a data transfer between a host and the storage system, controller cache will contain either all the old data or all the new data, not a mix of old and new data. This option has a slight performance cost because it maintains a secondary copy of data in cache so that if a data transfer is not completed, the old cache data can be restored. See also host, storage system.
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. See also leftover and vdisk.
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 SC and the MC in a controller module. CAPI is always enabled. See also controller module, MC, and SC.
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. See also WBEM.
CIM Query Language	See CQL.
CIMOM	Common Information Model Object Manager. A component in CIM that handles the interactions between management applications and providers. See also CIM.
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. See also vdisk.
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). See also controller module.

controller enclosure	An enclosure that contains one or two controller modules. See also controller module, enclosure
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 *A* and the lower one is designated *B*. See also controller enclosure, EC, FRU, host, MC, and SC.
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. See also CIM.
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. See also storage system.
Data Encryption Standard	See DES.
dedicated spare	A disk that is reserved for use by a specific vdisk to replace a failed disk. See compatible disk and vdisk.
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. See also map, explicit mapping, host, masking, and volume.
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. See also host.
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 expasion modules. Drive enclosures can be connected to a controller enclosure to provide additional storage capacity. See also controller enclosure, enclosure, and expansion module.
drive spin down	See DSD.
DSD	Drive spin down. A power-saving feature that monitors disk activity in the storage system and spins down inactive SAS disks based on user-selectable policies. See also storage system.
Dynamic Host Configuration Protocol	See DHCP.
dynamic spare	An available disk that is automatically assigned, if the dynamic spares option is enabled, to replace a failed disk in a vdisk. See available disk, compatible disk, and vdisk.
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 also controller module, EMP, expansion module, and SES.

EMP	Enclosure management processor. An EC subsystem that provides SES data such as temperature, PSU and fan status, and the presence or absence of disks. See also EC, PSU,. and SES.
enclosure	A physical storage device that contains disk drives and other FRUs. See also FRU.
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 *A* and the lower one is designated *B*. See also drive enclosure, EC, FRU, and host.
explicit mapping	Access settings for a host to a volume that override the volume's default. See also default mapping, host, and masking.
extrinsic methods	Methods which are particular to a specific class in SMI-S.
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 also Host, LUN, recovery, volume, and WWN.
FC-AL	Fibre Channel Arbitrated Loop. The FC topology in which devices are connected in a one-way loop. See also 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 to replace a failed disk. See also compatible disk and vdisk.
HBA	Host bus adapter. A device that facilitates I/O processing and physical connectivity between a host and the storage system. See also host and storage system.
host	An external port that the storage system is connected to. The external port may be a port in an I/O adapter in a server, or a port in a network switch. See also storage system
host port	A port on a controller module that interfaces to a host computer, either directly or through a network switch. See also controller module and host.
host bus adapter	See HBA.
initiator	See host.
I/O Manager	An MIB-specific term for a controller module. See also controller module and MIB.
I/O Module	See IOM.
IOM	I/O Module. An IOM can be either a controller module or an expansion module. See also controller module and expansion module.
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 disks'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. See also metadata and vdisk/
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. See also loop.
Logical Block Address	See LBA.
Logical Unit Number	See LUN.
loop	FC-AL topology. See also FC-AL.
Loop Initialization Primitive	See LIP.
LUN	Logical Unit Number. A number that identifies a mapped Volume to a host. See also host, map, and volume.
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 (or 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 also default mapping, explicit mapping, host, host port, initiator, and volume.
masking	A volume-mapping setting that specifies no access to that volume by hosts. See default mapping, explicit mappint, map (or mapping), and volume.
MC	Management controller. A processor located in a controller module that is responsible for human-computer interfaces and computer-computer interfaces, including the WBI, CLI, and FTP interfaces, and interacts with the SC. See also controller module, SC, and WBI.
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 ID, vdisk ownership, volumes in the vdisk, host mapping of volumes, and results of the last media scrub. See also host, map (or mapping), vdisk, and volume.
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 also host, map, and volume.
network port	An Ethernet port on a controller module through which its MC is connected to the network. See also controller module and MC.
network time protocol	See NTP.
NTP	Network time protocol.
object identifier	See OID.
OID	Object Identifier. A number assigned to devices in a network for identification purposes.
orphan data	See unwritable cache data.
Partner Firmware Upgrade	See PFU.
PFU	Partner Firmware Upgrade. The automatic update of the partner controller when the user updates firmware on one controller in a dual-controller system.

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	The FC 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.
PSU	Power Supply Unit. The power supply FRU. See also 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 also failover, host, LUN, volume, and WWN.
remote syslog support	See syslog.
SC	Storage Controller. A processor located in a controller module that is responsible for RAID controller functions. The SC is also referred to as the RAID controller. See also controller module.
SCSI Enclosure Services	See SES.
secure hash algorithm	See SHA.
secure shell	See SSH.
Secure Sockets Layer	See SSL.
Self-Monitoring Analysis and Reporting Technology	See SMART.
Service Location Protocol	See SLP.
SES	SCSI Enclosure Services. The protocol that allows the initiator to communicate with the enclosure using SCSI commands. See also enclosure and initiator.
SFCB	Small Footprint CIM Broker. See also CIM.
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. See also SNIA. The interpretation of CIM for storage. It provides a consistent definition and structure of data, using object-oriented techniques. See also CIM.
SNIA	Storage Networking Industry Association. An association regarding storage networking technology and applications.
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. See also controller enclosure and drive enclosure.
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 either controller's host ports. ULP incorporates ALUA extensions. See also host, host port, LUN, map, and volume.
Unified LUN Presentation	See ULP.
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. Unwritable cache data is also called orphan data. See also volume.
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. See also WBI.
vdisk	A virtual disk comprised of the capacity of one or more physical 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. See also host and vdisk.
WBEM	Web-Based Enterprise Management. A set of management and internet standard technologies developed to unify the management of enterprise computing environments. See also CIM.
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 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.
WWPN	World Wide Port Name. A globally unique 64-bit number that identifies a port.

Index

F

G

H

I

L

M

N

O

P

R

S

T

U

V

W